

A comparative survey of Personalised Information Retrieval and Adaptive Hypermedia techniques

Ben Steichen^{a,*}, Helen Ashman^b and Vincent Wade^a

^a Knowledge and Data Engineering Group, School of Computer Science and Statistics,
Trinity College Dublin, College Green, Dublin 2, Ireland
{Ben.Steichen, Vincent.Wade}@cs.tcd.ie

^b WebTech and Security Lab, School of Computer and Information Science,
University of South Australia, Mawson Lakes, SA 5095, Australia
helen.ashman@unisa.edu.au

Abstract. A key driver for next generation web information retrieval systems is becoming the degree to which a user's search and presentation experience is adapted to individual user properties and contexts of use. Over the past decades, two parallel threads of personalisation research have emerged, one originating in the document space in the area of Personalised Information Retrieval (PIR) and the other arising from the hypertext space in the field of Adaptive Hypermedia (AH).

PIR typically aims to bias search results towards more personally relevant information by modifying traditional document ranking algorithms. Such techniques tend to represent users with simplified personas (often based on historic interests), enabling the efficient calculation of personalised ranked lists. On the other hand, the field of Adaptive Hypermedia (AH) has addressed the challenge of biasing content retrieval and presentation by adapting towards multiple characteristics. These characteristics, more typically called personalisation "dimensions", include user goals or prior knowledge, enabling adaptive and personalised result compositions and navigations.

The question arises as to whether it is possible to provide a comparison of PIR and AH, where the respective strengths and limitations can be exposed, but also where potential complementary affordances can be identified. This survey investigates the key techniques and impacts in the use of PIR and AH technology in order to identify such affordances and limitations. In particular, the techniques are analysed by examining key activities in the retrieval process, namely (i) *query adaptation*, (ii) *adaptive retrieval* and (iii) *adaptive result composition and presentation*. In each of these areas, the survey identifies individual strengths and limitations. Following this comparison of techniques, the paper also illustrates an example of a potential synergy in a hybridised approach, where adaptation can be tailored in different aspects of PIR and AH systems. Moreover, the concerns resulting from interdependencies and the respective tradeoffs of techniques are discussed, along with potential future directions and remaining challenges.

Keywords: Adaptive Hypermedia; Personalised Information Retrieval; Query Adaptation; Adaptive Composition; Adaptive Presentation; Adaptive Web.

1. Introduction

Many information systems have attempted to solve the information overload problem that querying information seekers are currently facing. However, despite being very efficient, traditional Information Retrieval (IR) techniques often follow the one-size-fits-all paradigm by delivering the same information in the same form and order for every user with the same query. Since different user information needs and queries arise in varying contexts with different intentions (although possibly not expressed in the query), research has started to focus on not just *retrieving potentially relevant documents*, but rather *delivering*

* Corresponding author. Tel.: +353 1 896 1335; fax: +353 1 677 2204; e-mail address: Ben.Steichen@cs.tcd.ie

tailored, adapted and personalised information presentations (Dumais, 2009). This development has sparked off the notion of Personalised Information Retrieval (PIR), which attempts to modify and evolve established IR techniques in order to produce more personally relevant results. Such systems tend to represent users with simplified personas, which are often based on historic interests or user location properties (e.g. geographical location, language prevalent in a region). Initial evidence has emerged that some of these PIR techniques have been applied within popular web search engines, however little detail has been published so far. Although such statistical approaches enable the efficient calculation of personalised ranked lists, other considerations such as user context¹ or preferences are often neglected. Similarly, while user expertise has been recognised as having a significant influence on search behaviour, little attention has been paid on adapting towards such differences (Al-Maskarian and Sanderson (2011)).

By moving towards adaptation and personalisation, PIR increasingly faces challenges that have been addressed extensively by the field of Adaptive Hypermedia (AH). The approaches and techniques that have arisen in this field have been inherently focused on adapting information presentations to varying user needs, backgrounds and contexts in order to provide personalised responses. Such an integration of multiple considerations, also called personalisation “dimensions” (Wade 2009), allows the non-linear composition of results according to a user’s current goals, preferences and context. The non-linear approach of AH with respect to composition stems from its background in hypertext research, as opposed to the document-centric ranking paradigm of traditional IR systems. This notion of personalised result composition is therefore often missing in PIR systems, which typically only attempt to perform some form of result summarisation or extraction. However, there have been substantial shortcomings to AH techniques as well, most notably in terms of applicability to large sets of documents. Since most applications have been confined to small sets of handcrafted content, AH has so far not had the real-world impact that IR has enjoyed over the last decades through web search engines (Brusilovsky and Henze 2007). Thus it might be argued that AH needs scaling up in order to incorporate external documents, whereas PIR needs scaling down to an individual’s preferences and context.

This survey analyses the key affordances, techniques and impacts of PIR and AH, in order to find their respective strengths and weaknesses. The motivation for the comparison is the fact that both research fields aim to solve similar issues and challenges but approached from opposite directions: PIR provides personalised information predominantly through adaptive document ranking techniques, whereas AH delivers personalised information through adaptive compositions and presentations. There are many areas where one field could benefit from the other, and, more importantly, there is a vast potential for the fusion of selected techniques and approaches. This could potentially lead to *web-scale personalisation*, a goal that each of the fields are currently striving to achieve on their own. Web-scale personalisation can be defined as the application of personalisation on a just-in-time basis over existing, third party data. This personalisation should address a user’s particular context, goals, needs and preferences by composing information presentations from a variety of external data sources. By contrast, current PIR systems often apply large-scale personalisation on the sole dimension of previous interests, whereas AH systems apply multi-dimensional personalisation only over the data they manage.

There are some other personalisation research fields that have also grown in prominence in the last 5 years, e.g. recommender systems and social search systems. For recommender systems, there are three common categories: content-based recommender systems, collaborative recommender systems, as well as hybrid systems, which combine techniques from the former two (Burke, 2007). Content-based systems make use of the features associated with items as well as the interest rating a user has given to them. Such systems typically make use of statistical techniques commonly found in Personalised Information Retrieval systems, treating recommendations as a personalised classification problem (Burke, 2007). Collaborative filter based recommender systems are less content-oriented and users are typically presented with documents or items that are recommended by users with similar interests (e.g. AMAZON² recommendations). The

¹ Context, as defined in (Dey, 2001) “is any information that can be used to characterise the situation of an entity.” While such information can include data from physical sensors such as GPS, accelerometers or indoor-location tracking (e.g. in Brown and Jones (2001), Maekawa et al. (2009), Noh, et al. (2011)), the notion of context within this survey relates to the situational information need of a person, e.g. including the current task or the person’s domain knowledge.

² <http://www.amazon.com>

neighbourhood of such peer users is typically only calculated based on users' rating histories. Similarly, social search systems typically use such collaborative filtering techniques on user interest ratings in order to identify documents that were of interest to similar peer users. With the growth of online communities, these techniques might become increasingly powerful for future adaptive and personalised search.

While PIR and AH can also be regarded as performing such "recommendations" (e.g. recommending links), there are many other adaptive techniques and behaviours that these systems perform (e.g. query expansion, adaptive composition, adaptive presentation). This distinguishes the fields from traditional Recommender Systems, which mainly focus on recommending items through adaptive retrieval.

The focus of this survey lies in the techniques and technologies that are used specifically by PIR and AH (and which have less relationships with Recommender Systems). First of all, the reason for specifically analysing PIR and AH lies in the novelty of comparing and possibly combining the complementary affordances of the techniques particular to these fields. The authors believe that there is a strong but undiscovered synergy between the two fields, and that there is vast potential for combining respective techniques and technologies. Secondly, in order to keep the desired depth in PIR and AH without weakening the focus of the paper, this survey particularly concentrates on the analysis on these two fields. However, while it is out of the scope of this survey to also include Recommender Systems, it would be interesting to have another survey on Recommender Systems and PIR, as well as Recommender Systems and AH. We refer the interested reader to Pazzani and Billsus (2007) for a survey of content-based recommender systems. Similarly, while social techniques are becoming increasingly popular for recommending information (by using social features for an adaptive retrieval process), such systems are not the focus of this survey. However, we refer the interested reader to Schafer et al. (2007) for a survey of such techniques.

For the purposes of analysing where adaptation can influence the search process, it is possible to characterise search technologies into three general stages: *query adaptation*, *retrieval adaptation* and *result composition/presentation*. First of all, the *query adaptation* step entails the analysis of an original query and its adaptation based on the available user information. Secondly, the *retrieval algorithm* can be *adapted* to incorporate personalisation features that tailor the retrieval to provide results that are deemed more relevant to the specific querying user. After the adaptive retrieval has produced personalised results, this information can be *adapted* further by *composing and presenting* the results in a coherent and personalised manner. Since the retrieval stage mostly provides unstructured results, these results need to be composed into a structured form that is most suitable for the querying user. This composition ideally guides the user through the results and possibly provides a structure that helps the user explore the information in the most effective and efficient way. In addition to the result composition, the actual presentation can be adapted to suit the particular user.

The remainder of this survey is structured as follows. Section 2 provides a comparison of the information access paradigms underlying the fields of AH and PIR. This is followed by an analysis of each of the retrieval process stages, namely *query adaptation* (section 3), *adaptive retrieval* (section 4) and *adaptive composition and presentation* (section 5). The key approaches and techniques are identified and presented using examples from the research literature. Moreover, the different approaches are compared and critiqued broadly across their objectives, models, algorithms and scalability. The overall findings are presented in section 6, and the potential for a hybridisation approach is outlined in section 7. This approach is presented along with an example scenario, where such selective hybrid techniques could enable improved adaptation and personalisation. During the investigation of the hybridisation approach, the challenges resulting from interdependencies and the respective tradeoffs of the various approaches are also examined, beginning to identify when particular parts of the retrieval process need to be influenced. Finally the paper concludes by outlining several future directions (section 8).

2. Information Access Paradigms of PIR and AH

One of the major distinctions between PIR and AH lies in the overall information access paradigms underlying both areas. This distinction has driven most of the research in both fields and has led to their respective techniques and technologies towards adaptation and personalisation.

PIR is based on the standard Information Retrieval model, which is traditionally focused on the retrieval of documents that are relevant to a unitary query. While PIR extends this model by taking into account

historical interactions, the paradigm remains *finding the most relevant documents* for a single user query. This fundamental underpinning of PIR makes such systems particularly suitable for the general information access paradigm of *searching by query*, where it is assumed that a user can express their information need in a relatively precise user query. However, this assumption is not always correct, especially in cases where users are uncertain about their actual information needs or the correct query terms to express these needs.

AH systems stem from the information access paradigm of *searching by browsing*, where users generally have less precise information needs and therefore need to surf and explore pages. AH is facilitating this type of search by *providing the most relevant browsing content and links* with respect to a rich representation of user characteristics (such as preferences, history or prior knowledge). AH is assisting a user's information exploration by creating or adapting the navigation across hypertext and hypermedia using e.g. link creation, link/content hiding or link/content annotation.

This distinction in information access paradigms has also driven the application of the respective techniques towards specific user information needs. In (Jansen et al., 2008), search needs are classified into three categories, *Navigational* searching, *Transactional* searching and *Informational* searching. *Navigational* searches typically only require one precise answer, for example the web address of a particular company or the homepage of an individual. Similarly, *Transactional* searches are very specific, as the main goal is to purchase a particular product or use a particular service. As PIR focuses on finding the most relevant documents for precise user queries, it is particularly suited for these information needs. However, these two types of searches are shown to constitute only 20% of current searches on the web (Jansen et al., 2008). The remaining 80% of searches can be classified as *Informational* searching, i.e. a user searching for deeper information on a particular topic. Current (P)IR solutions typically do not assist this *Informational* searching adequately, as often more than a few very precise results are needed to fill a user's knowledge gap. AH systems have inherently focussed on providing assistance towards such *Informational* searching needs, since they typically assist users in gaining deep knowledge in a particular domain. In particular, AH systems have proven successful in e-learning systems, where a user's information need is to learn about a particular topic. This represents a good example of a search reaching beyond a few precise documents and corresponds to a case where a user is uncertain about the range of query terms to choose. However, while AH systems have proven effective for such "formal" learning tasks, their closed corpus nature has prevented a more widespread applicability for web-scale *Informational* searching (i.e. "informal" learning activities).

Despite PIR and AH stemming from distinct information access paradigms, both fields effectively strive towards similar aims of providing the most *personally relevant information* to an individual user. This has led to many conceptual commonalities between the techniques of both fields, which can potentially complement each other in order to better assist a user's *information exploration* on a web-scale. The aim of assisting such searches is also actively researched in the area of *Exploratory Search* (Marchionini, 2006). This field distinguishes itself from standard search through the following 6 characteristics (White and Roth, 2009): 1) multiple query iterations, 2) open-ended information needs, 3) close coupling with task context, 4) combination of browsing and focused searching, 5) possible collaboration of multiple people and 6) advanced system evaluation with respect to learning, insight, task outcomes and system utility. One could thus argue that a hybrid of both PIR and AH systems could potentially be characterised by this search paradigm. By analysing and contrasting the various PIR and AH techniques, this survey is able to provide a set of affordances that such a hybrid system's components require in order to improve user assistance across *Informational* and *Exploratory* search.

3. Query adaptation

The adaptation of user queries has been explored from many angles in several distinct research fields. The classic approach to the elicitation of an information query consists of a short interaction involving a user specification of a set of keywords. Although in theory these terms denote the user's original informational need, they are often short (Spink and Jansen, 2004) and possibly ambiguous or incomplete due to common natural language problems such as homonymy and polysemy (Krovetz, 1997). This is reinforced by users generally providing low commitment to search interactions and having overly high expectations with respect to the search system (Jansen, Spink and Saracevic, 2000). Since this is the main information entering an information retrieval engine, several attempts have been made to enhance the original query to more accurately reflect the user's perceived intention.

In a survey about semantic-based approaches, Mangold (2007) provides an overview of the different types of adaptation that can be applied to user queries. It is argued that most adaptation techniques are used to overcome the problem of ambiguity. The types of adaptations identified are: manual query modification (a user reformulating a query), graph-based modification (a user navigating a graph, which in turn generates new queries), query augmentation (expansion), query trimming and query substitution. It is argued that different techniques may perform better for different purposes (e.g. query expansion increases precision whereas query trimming increases recall), with some systems now using a combination of different techniques. The survey below provides a summary of the implemented approaches and techniques.

Classical IR solutions for query adaptation have mainly consisted of statistical query expansion techniques, with the earliest systems performing Relevance Feedback (RF) from explicit user relevance judgments (Rocchio, 1971). In this approach, a user can specify which documents are relevant after an initial search query has produced a set of results. The most frequently occurring terms from the initially relevant documents are then used to expand the original query. This technique, albeit increasing precision and recall in research evaluations, is compromised by the fact that users are required to invest extra effort compared to regular search (since users have to go through multiple iterations to refine their query). Given that most users put very little effort into search, this could be considered as too cumbersome. A technique called Pseudo-Relevance Feedback (Xu, 1996) has automated this process by automatically selecting the top ranked documents as relevant information sources to perform relevance feedback. No user intervention is required, although it is arguable if this technique can still be regarded as personalised adaptation. Another popular statistical method for choosing expansion terms uses the concept of co-occurrence. For a certain term x in a user query, terms that frequently co-occur with x in the document base are shown to be excellent query expansion candidates (Kim and Choi, 1999). This technique is also evident in web search engines in the form of query suggestions provided either while eliciting the query or after an initial search has been conducted.

More recently, researchers have focused on improving the *personalisation in relevance feedback* by making use of users' search history and context. Koutrika and Ioannidis (2005) propose to mine search queries and relevance judgements into a user profile. This information can then be used to perform query disambiguation using query expansion based on the mined relationships. Logical operators such as AND, OR, NOT can be used to issue a new personalised query for the particular user. In Pitkow et al. (2002), a user bootstraps a user profile by importing bookmarks into the *Outride* tool. The bookmark links are then mapped to the top 1000 ODP³ categories in order to create a weighted user profile. Additionally, a recent account of user interests is kept by continually updating the profile with a user's click through history. When issuing a new query, the terms are compared against the user profile in order to select appropriate expansion terms. The mapping of websites to ODP categories is achieved by first training a keyword-based classifier on a set of documents listed for each category (see Figure 1).

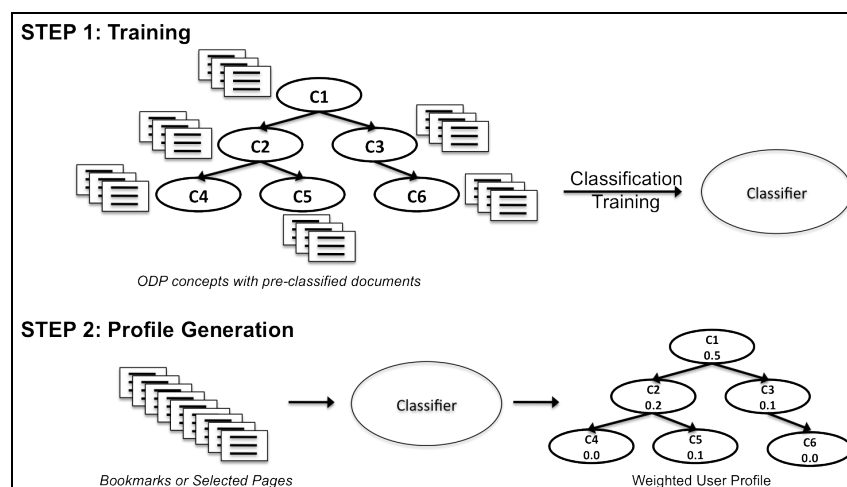


Figure 1. User profile generation using ODP categories

³ The Open Directory Project (ODP). <http://dmoz.org>

The user profile is then generated by classifying the bookmarked websites (i.e. the textual content) to ODP categories. This classification may also result in weighted similarity scores, allowing a more fine-grained representation of user interests (Micarelli, et al., 2006). Similarly, in Teevan, et al. (2005) the corpus for relevance feedback is not chosen from top ranked documents of an initial query, but from a user profile. This information is gathered from users' personal desktop applications (e-mail, word processing, etc.), as well as their browsing and search histories. It is argued that this information better denotes the short and long term user preferences. A similar approach was also taken by Chirita, et al. (2007), who were using a Personal Information Repository to choose expansion terms. Additionally, the number of query terms is adaptively chosen based on the query clarity. This query feature is calculated using both scope of the query (i.e. its document coverage), as well as the query language model diversion from the document collection language model. Experiments show that further improvements can be achieved when adding such adaptive functionality.

Another form of implicitly derived relevance feedback mines users' interactions with search engines or other information systems to promote or reorder results returned to users. The outcomes are purely statistical and do not require any content-specific labelling or processing of items, and also do not require any explicit statement of interest from the user – the relevance feedback is purely a byproduct of the user's other activities. An example of such implicit relevance feedback occurs in the use of clickthrough data and coselection data – a click (or clickthrough) is where a user selects a specific result, which is interpreted as positive feedback on the relevance of the selected result to the search term, and can thus be used for result reordering. A coselection is where two or more results are selected from the same search, which implies not just relevance of the selected results to the search term, but additionally implies relevance of the selected results to each other. When aggregated, it is possible to use coselections to disambiguate the search term (Truran et al., 2005).

In Calegari and Pasi (2008), it is proposed to create a fuzzy ontology from users' search query histories and the documents stored on their personal workstations. Relationships are represented as simple numerical relationship strengths, which are displayed visually when a user issues a new query. The user may then choose related terms in order to expand the original query. Similar ideas are explored in (Ahn et al., 2010), where named entities are extracted from an initial set of retrieval results. These entities are then organised by their prominence and displayed along with the result list. The user can then select one or more named entities as expansions to the original query terms. Figure 2 gives an example of this interface, where the user has selected “Franz Schausberger” (left) and “Salzburg” (right) in addition to the original user query “train fire”. Such efforts move closer to ideas explored in Adaptive Hypermedia, which inherently make user of semantically rich conceptual models of user interests.

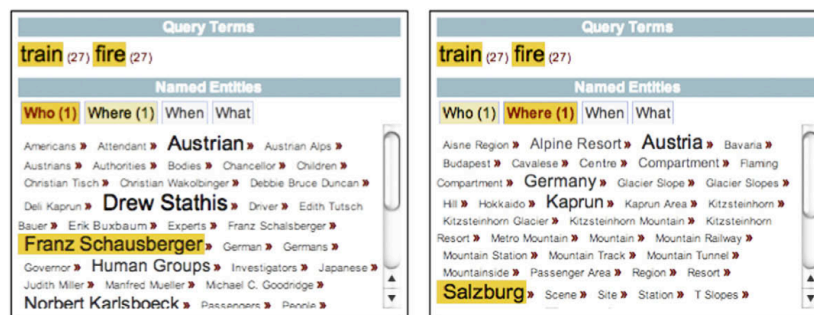


Figure 2. User interface for manual selection of additional query terms (Ahn et al., 2010)

Although the field of Adaptive Hypermedia typically does not explicitly involve user queries (as it has historically been based on browsing as opposed to searching), several related ideas can be found in the use of semantic retrieval techniques. Just as Adaptive Hypermedia systems use metadata, concepts and conceptual relationships to drive their adaptation, semantic-based retrieval makes use of such metadata and rich conceptual models to adapt the retrieval process (see section 4.2.). In addition to this retrieval adaptation, several semantic-based systems also make use of such models to produce initial query adaptation. A survey by Bhogal, Macfarlane and Smith (2007) reviews the use of ontologies for query expansion and argues that corpus-independent knowledge models can be used to provide disambiguation

using the implicit semantic knowledge embedded in such models. Many surveyed systems make use of the Wordnet ontology synsets (Miller, 1995). This general ontology can be used to find related concepts, such as subclasses and superclasses, as well as synonyms to perform query disambiguation. Navigli and Velardi (2003) make use of ontologies in several ways by using a multitude of semantic relationships, e.g. hyperonymy (is-a), meronymy (has-a), similarity, etc. They compare the effects of several expansion techniques using general knowledge bases and conclude that a general improvement occurs over unexpanded queries. They argue that shorter queries benefit more from expansion due to the ambiguity associated with them. They also propose emergent semantic similarities between concepts by searching for “common nodes”. This technique chooses words for expansion based on the fact that they have similar synsets to the original query terms. The similarity is calculated by searching for nodes that both the original and the expansion node have in common.

Rocha, Schwabe and Aragao (2004) take an even more sophisticated approach by applying spreading activation to the underlying knowledge base. A query is expanded by a set of terms that have been “activated” in the knowledge graph due to their semantic relationships with the initial terms. This activation spreading is using ontology relationships, coupled with different weights attributed to each of these relationships. Depending on the particular domain or user preferences, stronger weights could be attributed to particular relationships in order to provide a “personalised” activation set. An important ontological relationship that is only rarely explored is antonymy (opposite of). Burton-Jones, Storey and Sugumaran (2003) argue that by including this in an expanded query with a preceding logical negation, query disambiguation can be enhanced further than by simple hyperonymy-based solutions. An even richer use of the semantic reasoning capabilities in ontologies is proposed by Linckels, Sack and Meinel (2007), where user queries are translated into Description Logic (DL). Although their work concentrates on searching over a closed knowledge base, it is worth noting that the semantic power of DL in ontologies could be used in several ways to reason about the expansion/reduction of concepts, as well as the logical operators separating the different terms.

It is interesting to note that the above methods all assume that an initial user query is adapted to form only one new query. However, Radlinski and Dumais (2006) propose to generate a set of related queries from an initial query by using a large sample of query logs. This set of queries is then used to generate results that can be reranked during the composition stage. The query choice is calculated by analysing the probabilities of which queries from the log are most likely to follow the initial query based on past usage patterns. Similarly, Glover et al. (2001) generate a set of modified queries from the initial query in their Inquirus2 system and then submit these to a selection of search engines (see Figure 3).

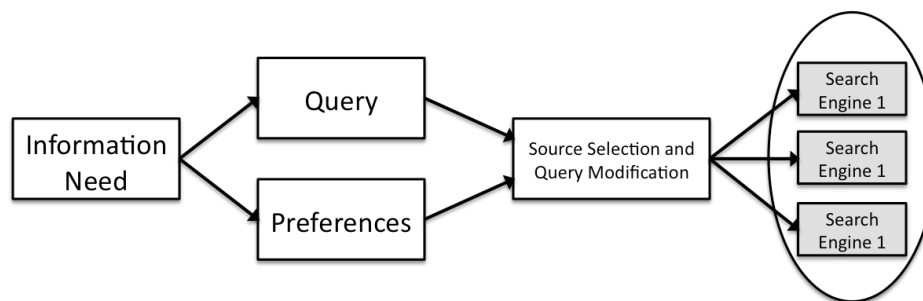


Figure 3. Personalised meta-search engine (Glover et al., 2001)

The modified queries however do not take into account semantics related to the query, as they rather attempt to broaden the search by appending terms such as “links” or “resources” for particular search types. The type of words added to the queries depends on a “need category” specified by a user, which denotes a certain query intent. Additionally, the type of search engine chosen, as well as particular query constraints (such as recency) are chosen with respect to the particular need category. Similarly, in Kumaran and Allan (2008), the authors compare Interactive Query Reduction (IQR), Interactive Query Expansion (IQE) and a hybrid Selective Interactive Reduction and Expansion (SIRE). SIRE selects the top five sub-queries and the top five expansion queries and presents the results to the user, who can then guide the system implicitly to preferred queries. Results show that SIRE outperforms a baseline system as well as both IQR and IQE.

3.1 Summary and Critique

The analysis of query adaptation and personalisation shows a broad range of techniques that have been developed using either statistical or semantic approaches. A summary of these techniques can be found in Table 1.

Table 1. Summary of Query Adaptation Techniques

	Classic Statistical	Personalised Relevance Feedback	Semantic Techniques	Meta-Search
Objectives	Query Disambiguation/ Personalisation	Query Disambiguation/ Personalisation	Query Disambiguation	Result Diversification
Knowledge Base	Document Set	Document Set, Personal Information Repository (PIR)	Semantic Knowledge Base	Document Set
Document Scale	Large Open Corpus	Large Open Corpus	Small to Medium Closed Corpus	Large Open Corpus
Model Requirements	none	ODP categories	General/Domain Ontology	Diversification Strategy
User Involvement	Relevance Judgements	Relevance Judgements	Explicit User Feedback	Relevance Feedback, User intent elicitation
User Model	none	PIR Relationships, Weighted ODP categories, Named Entities	none	none
Adaptation algorithm	Statistical Similarity (Keyword-based)	Statistical Similarity (Keyword-based)	Semantic Reasoning	Statistical/ Strategy- driven
Query Modification	Weighted Query Expansion	Query & Logical Operator Expansion/ Trimming	Query Expansion/ Substitution	Query Substitution/ Expansion/ Trimming, Multiple Query Generation

Classic Statistical techniques are shown to focus on bulk document analysis and keyword similarity in order to produce expanded queries. Since the focus lies on large document collections, no additional knowledge bases or user modeling components are required to apply these techniques. It can therefore be applied to large open corpus domains with high efficiency. This constitutes the biggest advantage of this type of system and it is therefore the most widely deployed approach on a web scale.

Although users can perform relevance feedback on initial sets of ranked documents, these systems do not keep a persistent record of particular user preferences, nor do they attempt to model the document space in a structured way. The personalisation aspect of these techniques is therefore limited, as user queries are regarded as *ad hoc* interactions. The relevance feedback never gets aggregated into some larger comprehension of relevance to the search term and thus each user needs to provide feedback for the same query. Moreover, these techniques only consider query expansion, not taking into account that query reduction or substitution might be more effective in certain cases.

Personalised Relevance Feedback techniques have attempted to acquire additional personal information of users in order to perform improved statistical similarity measures. Furthermore, not just query expansion

but also query trimming and logical operator additions are considered for the adaptation of queries. Since the main focus still lies on statistical similarity measures, these systems can also be applied to large open corpus document bases. Personal information repositories as well as ODP categories are used as simplified knowledge bases that improve the personalisation effectiveness compared to classical statistical models. Again, the most important advantage of these systems is their current applicability on a web scale, since the knowledge base creation and document classification make use of efficient statistical similarity measures.

However, user models are still represented as simple keyword or keyword-relationship vectors, which contain little semantics in order to infer personalised query adaptation strategies. Moreover, these techniques do not perform semantic domain model reasoning, nor do they contain strategies for result diversification.

Semantic Techniques introduce the notion of a semantic domain model, which incorporates rich relationship information as well as reasoning rules. Query disambiguation is achieved using these relationships in conjunction with explicit user feedback.

However, the notion of a user model has been widely unexplored by these systems, leaving the adaptation process on a non-personalised level. Additionally, queries are not diversified, as the main objective of these systems is to simply disambiguate a single query. Furthermore, semantic-based techniques have largely been confined to closed corpus domains due to the current reliance on semantically rich models. Since it is difficult to create these models automatically, most systems have been built manually in order to query a small to medium sized digital library.

A solution to the non-adaptive nature of semantic retrieval techniques would be to use more advanced Adaptive Hypermedia techniques such as user modelling and personalised strategies in order to better adapt queries. However, the scalability issue described above would still persist, as both fields suffer equally from high model and metadata requirement costs.

Several *Meta Search* systems have explored the diversification of search results by rewriting/generating sets of multiple queries, using either statistical similarity measures or particular expansion strategies. The term ‘meta search’ is used here to denote both systems that generate multiple queries to the same underlying search engine and systems that send queries to several different engines. Since all these systems do not just focus on providing improved single queries, they capture a greater breadth of search results. Additionally, by relying on statistical similarity measures, these systems are able to operate on large open corpora. Personalisation occurs after a user provides relevance feedback or an intent elicitation. Using this information, systems can adapt the resulting queries more precisely towards particular information needs.

However, the notion of expansion strategies introduced in Glover et al. (2001) has been confined to simple hardcoded rules that are applied for every user with the same query. Firstly, this leads to a scalability issue for the creation of new rules, since they each have to be created manually. Secondly, they have yet to be personalised in order to adapt the strategy to the particular querying user. Furthermore, diversification strategies generally do not make use of a knowledge base or semantic user model in order to reason about adaptive query expansion/trimming.

3.2. Comparison across Query Adaptation techniques

In conclusion, the analysis of query adaptation approaches reveals a variety of techniques that have been applied in order to improve initial user keyword queries.

The *objective* of most techniques has been confined to simply disambiguating user queries, based on the assumption that the sparse set of initial keywords does not contain enough information to provide precise results. Furthermore, apart from several personalised relevance feedback systems, most approaches apply adaptation on a non-personalised level, largely ignoring user intent identification and processing. Additionally, only a few systems move away from the generation of single expanded queries towards query diversification using multiple queries. A clear lack can be identified in the exploration of combining the different objectives, which would consist of an adaptive strategy for choosing appropriate and personalised adaptation techniques.

In terms of *knowledge base and model requirements*, there is a clear distinction between statistical and semantic techniques. Statistical solutions generally only require a large document set in order to calculate

(personalised) relevance scores. Recent systems make use of personal information repositories and ODP categories in order to provide additional result classification and personalisation. This independence from complex domain models constitutes a great advantage over semantic solutions, which can only be applied to domains where such models have been created. Although the full combination of both scalability and domain model reasoning would constitute the ideal solution, a compromise could be reached by using statistical document set analysis together with semi-formal domain ontologies.

As mentioned before, *scalability* has been the main focus of statistical IR systems. This constitutes a clear advantage over semantic retrieval systems, as it allows them to be applied to any large open corpus. The challenge of such systems hence lies in improving semantic query adaptation techniques in order to improve their current scalability issues. The incorporation of statistical similarity measures into semantic based retrieval systems as well as Adaptive Hypermedia systems proves necessary, as the loosening of metadata requirements could provide more scalable, lightweight solutions.

In order to perform personalisation, systems either require *User Involvement* such as *ad hoc* relevance judgements or they create *User Models* in the form of more complex weighted sets of keyword relationships and category classifications. Users are often required to manually judge sets of relevant documents or they are asked to pick expansion terms based on the mined relationships. The more semantics that are encoded in the user profile, the more is it possible for systems to infer personalised query adaptation strategies. Personalised relevance feedback systems have been successful in mining large amounts of basic information from users' search histories and personal workstations. An improved solution would index this information using increased semantic relationships in order to represent users' preferences and requirements more precisely. Additionally, this would allow a system to apply personalised adaptation strategies that could infer the right adaptation techniques for the particular user and context. By contrast, the sole dimension of adaptation that is currently being used is user interests, mined from previous interactions. Especially the user intent of a query would be of great use to a system in order to create a customised strategy for search result diversification. Currently, only the system provided by Glover et al. (2001) provides a basic implementation of this feature by requiring a user to explicitly choose a query intent.

Adaptation algorithms have been clearly divided into statistical similarity measures and semantic reasoning techniques. The clear advantage of statistical approaches lies again in the large scale applicability. However, it would be of great advantage if this strength is combined with domain and user model reasoning. Both methods do not exclude each other, they could rather be regarded as complementary. Especially in the case of insufficient semantic knowledge being available, statistical techniques could provide basic query adaptation functionalities.

Traditionally, the actual *query modification* focused solely on performing (weighted) query expansion. However, research has since been addressing a variety of adaptation techniques, each of which perform improvements over unmodified queries depending on the user and the context. An improved solution would adaptively choose the right modification based on the current user and domain model state. For example, a query could be both trimmed and expanded for a novice user in order to provide more guidance about the general domain. An initial adaptive algorithm has been presented by Chirita et al. (2007) where the number of expansion terms is chosen depending on the ambiguity of a query. This technique could be expanded by not only reasoning about the amount of keywords, but also about the type of modification that should be applied.

4. Retrieval adaptation

After an initial user query has been issued and possibly adapted, a retrieval engine is responsible for the retrieval of appropriate content. An adaptive personalised retrieval system targets the specific user information need and adapts to particular user preferences and context. Two main categories of techniques can be distinguished depending on the retrieval being based on statistical methods or on metadata-based algorithms.

4.1. Statistical methods

This category of algorithms is typically concerned with processing high volumes of data. Rather than replacing established IR/Web search algorithms, adaptive/personalised components are often attached as slight modifications or simply combined with the original search results. Ranked list scores are the preferred output of this type of algorithms and therefore the main focus of adaptive IR retrieval lies on improving the rank of documents that are relevant to the particular querying user.

An example of a modified retrieval algorithm is presented in Tanudjaja and Mui (2002). The ODP web taxonomy is used to capture user preferences, which then influence the HITS algorithm. HITS in its original form estimates the authority and hub values of a page solely using the given link structures (Kleinberg, 1999). However, the modified version gives more weight to pages that are related to positively rated pages (based on relevance feedback in the user's previous searches). Since each page can be mapped to the ODP taxonomy, relationships can be established by searching up and down the ODP tree for nodes that were previously explored. Similarly, Haveliwala (2003) provides a modified version of the popular link analysis algorithm PageRank. PageRank makes use of link information in order to provide a measure of popularity and authority of a page within a given set (Brin and Page, 1998). The presented modified version precomputes a set of topic-sensitive PageRank vectors for 16 ODP categories. At query time, the system first calculates the similarities of the query to the topics. Using these similarities, the system then adaptively calculates a linear combination of the topic-sensitive vectors for result ranking (see Fig. 4).




Figure 4. Topic-sensitive PageRank (Haveliwala, 2003)

Furthermore, users' past searches can be analysed to disambiguate a query, since the terms might be matched to several categories. Additionally, the authors propose to use the context in which the query was issued. For example by highlighting the search term on a website, the user provides valuable context, namely the complete web page where the term was chosen from. Another example of modifying an established retrieval algorithm is shown in Teevan et al. (2005). The well-known ranking technique of BM25 is modified in order to incorporate user interests. BM25 is a probabilistic ranking function that includes document and query term weights and which incorporates relevance feedback information (Croft et al., 2009). The proposed modified version performs a new type of personalised Relevance Feedback, with the user information being gathered from rich personal Desktop information. Therefore, it is able to infer relevance more accurately, since the re-ranking is based on documents that a user has actively interacted with.

In contrast to modifying established retrieval algorithms, many techniques calculate personalised result scores, which are then combined with original retrieval scores in order to determine the final result ranking. For example, the *Wifs* system in Micarelli and Sciarrone (2004) reranks initial search results (from Altavista) using similarity calculations between a user model and the returned documents. The user model is constructed using relevance feedback and contains terms that occur in the favourably rated documents. More specifically, terms that also occur in a manually constructed Terms Data Base (TDB) are considered as user model topics, whereas co-occurring words (found in the document but not in the TDB) are connected to the topics as co-keywords. Similarly, documents are represented using the occurring terms from the TDB, as well as non-TDB words co-occurring in the document. Several relevance calculations are then applied between the user model and the set of analysed documents in order to rerank the original results list.

Similar to the modified ranking algorithms, many combined ranking techniques make use of directory structures in order to rerank initial results. For example in Speretta and Gauch (2005), user profiles are constructed by mapping past queries and selected documents to ODP categories. The results for a newly issued query are analysed and mapped similarly to ODP categories and a similarity score is generated between the result documents and the user profile. This score is generated by multiplying the relative

weights of concepts in the result documents and the user profile. The higher this similarity score, the more the results are deemed to be personally relevant to the user. By contrast, Daoud et al. (2010) use graph-based ranking models, whereby the similarity score is calculated as a combination of the minimum common supergraph and the maximum common subgraph of the result documents and the user profile (similarly based on ODP categories). By combining the similarity score with the original search rank, result documents that are more relevant for the particular user consequently appear higher up in the list. A similar approach is taken in Pitkow, et al. (2002), where an original result list is re-ranked based on a user profile. However, this profile has been created by categorising bookmarking links (imported into Internet Explorer by the user) onto 1000 ODP categories, as well as categorising over time the search results selected by the user. By contrast, Xiang et al. (2010) only use the immediate context of a query, i.e. only successive queries (and their associated categories) are considered to be related (and consequently used for a personalised/contextualised score).

In addition to creating a weighted concept hierarchy, Liu, et al. (2004) associate different weights between particular keywords and the detected categories. For example, if a user interested in both “cooking” and “computers” has previously issued a query “apple” to retrieve “cooking” related documents, but not to retrieve “computer” related documents, the user profile will have a higher weight for “apple” in the “cooking” category. Therefore, this system has a higher degree of granularity over the system by Speretta and Gauch, since it takes into account more refined user preferences. An even more sophisticated classification is proposed in Stamou and Ntoulas (2009), where a topical ontology is created using ODP categories in conjunction with the Wordnet and SUMO ontologies. Users’ past queries are mapped to categories using several methods, including ontology traversal. The discovered topics are then used as a user profile during the rank combination phase in order to create a personalised document ranking.

4.2. Metadata-based approaches

The second type of adaptive retrieval systems heavily relies on rich models and document annotations. Semantic-based retrieval systems as well as Adaptive Hypermedia systems are built on the power of semantic knowledge engineering, aiming to achieve the vision of the Semantic Web (Berners-Lee et al., 2001). In terms of retrieval, techniques and algorithms utilised in Adaptive Hypermedia have been inherently conceived to provide adaptive result retrieval (in addition to the adaptive navigation and presentation functionalities presented in section 5.2). Since many prototype systems have been developed for the field of E-learning, most of the research has focussed on the personalised retrieval of learning content. However, the proposed techniques and algorithms are not exclusively applicable to this particular domain.

The notion of domain and user models has been proposed by the earliest *Adaptive Hypermedia* prototype systems, such as Interbook (Brusilovsky, et al., 2004) and AHA! (De Bra, et al., 2003). The domain model typically represents a conceptual view of the underlying domain, containing information about concept hierarchies, attributes and relationships. This model is independent from the underlying content and represents a more high-level model of the subject domain. For example, in an e-learning domain this model may contain various high-level topics to be covered within a subject, including relationships such as prerequisite requirements between topics. Interbook proposes the idea of mapping each document to a set of outcomes and prerequisites. It is argued that storing adaptation-specific information in external models assists the adaptive retrieval by allowing reasoning engines to infer which documents are relevant for the user. For example, the set of outcomes and prerequisites of a document allow a reasoning system to provide a user with a set of documents that should be visited before and after the current document. Additionally, using an overlay user model, Interbook is able to infer a student’s knowledge state for each of the domain concepts, which allows a personalised delivery/omission of information. Furthermore, the notion of a learning goal is proposed, which defines a particular sequence of documents in order to guide a student through the material. Similarly, the content adaptation in AHA! bases the inclusion/exclusion of fragments to be shown to the user on the state of the user model in relation to the domain model. Depending on previous user visits to certain pages, AHA! checks the suitability of a particular fragment in order to provide a student with personalised pages. In an improved version, called GALE (Smits and De Bra, 2011), it is shown that domain concepts can be distributed over several servers in order to decentralise the adaptive functionality. In Aroyo et al. (2004), it is proposed to incorporate external fragments into the AHA! system by making use of Information Retrieval (IR) results. The authors make the assumption that a resolver

ontology describing the search space can be mapped directly to the domain ontology. By doing so, the standard fragment inclusion/exclusion techniques can be performed on the domain concept level as in the case for the regular AHA! system. However, no experimental prototypes have been developed to evaluate these possibilities.

In the KBS hyperbook system (Henze and Nejd, 2001), documents are again linked to external models. However, as opposed to Interbook, the KBS hyperbook documents are only indexed with domain concepts, with the inter-concept relationships being defined in a Knowledge Model. Again, a user model is used to capture the actual knowledge of a student in order to compare this to the knowledge required to understand the topic in question. The system adaptively retrieves the set of concepts that the user should learn about first (the prerequisites). During the adaptive retrieval, all Knowledge Items that should be learned by the user are marked based on users' knowledge of the items and their prerequisites. The actual documents are then selected based on the document-concept indices that have been created *a priori*. All adaptation therefore occurs during the concept retrieval stage, rather than the content retrieval stage. It has to be noted that any document can be incorporated into the KBS hyperbook system as long as it is indexed to the concept space. The authors argue that a drawback of the system is that the learning dependencies are explicitly encoded into the concept space, although these relationships might differ for different scenarios. Depending on the intention of the knowledge engineer who designs the domain model, different assumptions might get encoded into the knowledge model. For example, different instructors might teach a certain subject using a different teaching strategy, leading to different prerequisite requirements encoded into the knowledge model. The authors acknowledge that this goes against the idea that a knowledge model should be independent of the particular teaching strategy.

The idea of selecting learning objects according to a particular strategy is shown in Farrell, Liburd and Thomas (2004). Initial XML search results are mapped to concept domain topics, for which further learning objects are retrieved in order to form a coherent learning path. The adaptive retrieval is based on the original query and the statistics collected for each topic during a mapping stage. Moreover, users can indicate their desired course duration depending on the time that is available to them. A drawback of the system is the fact that the particular retrieval strategy is encoded into the system rather than separated from the adaptive engine. In Conlan, et al. (2002) and Conlan and Wade (2004), a multi-modal, metadata-driven approach is proposed in order to provide this separation of concerns. Most importantly, the approach introduces an additional model called the narrative model, which encodes a set of generic strategies for presenting concepts. For example, this model can encapsulate an expert's knowledge of a domain and therefore provide guidance through appropriate course material. The implemented APeLS system executes the narrative by consolidating models in an Adaptive Engine (see Figure 5).

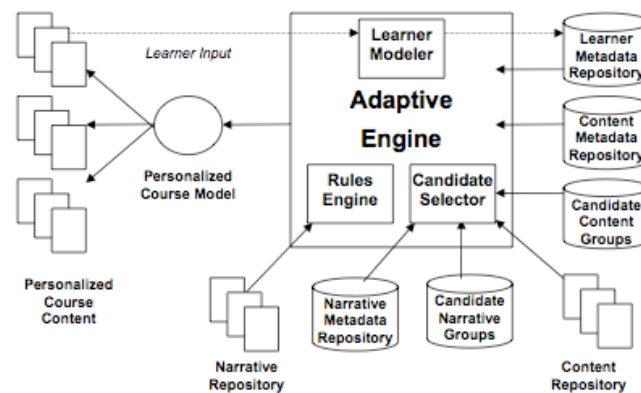


Figure 5. Architecture of the APeLS Adaptive Hypermedia Service Architecture (Conlan, et al., 2002) and Conlan and Wade (2004)

In terms of retrieval adaptation, the narrative contains the rules for which concepts should be selected, how they should be sequenced and which candidate content group should be considered for content retrieval. All adaptation hence occurs on the concept-level, enabling the narrative to be independent from the actual content. The adaptation rules can be based on the information that is available in any of the models, for

example user prior knowledge, media type preferences, history, cognitive style, etc. (each held in the Learner Metadata Repository). Narratives can be implemented using several technologies, such as rule-based (e.g. Drools⁴) or script-based languages (e.g. Javascript). Since all adaptation occurs on the concept level, the closed corpus could be expanded by taking a similar approach to the KBS hyperbook system. Furthermore, the idea of a narrative as a way to guide a user could be applied to different retrieval adaptation techniques in order to ensure coherent strategies.

The use of ontologies as concept domains is proposed in *semantic-based* systems such as the Personal Reader Framework (Henze, 2005). In this system, recommendations for learning resources are again based on the current learning progress of the user. Standardised metadata annotations are used in order to infer which learning object should be recommended/retrieved. Also, an alternative recommendation service is proposed, which could be based on the keywords that describe the objectives of the learning object in an ontology. This would allow different course materials to be used in the system. It is therefore argued that different recommendation services might be suitable in different situations. Furthermore, the use of more generalised ontologies for user observations and for adaptation makes it possible to share models among different applications. In Dolog, et al. (2003), it is proposed to fully move adaptive systems towards Semantic Web technologies in order to enable an adaptive Semantic Web. Due to the standardised formats, interoperability between applications would hence be greatly facilitated. It is argued that adaptive retrieval can be based on common logic-based languages, such as First-Order Logic. Since Semantic Web technologies are inherently focussed on reasoning capabilities, rule-based languages such as TRIPLE Sintek and Decker (2002) can be employed to reason across distributed metadata. The enhanced expressiveness of RDF⁵ and OWL⁶ allow the creation of comprehensive domain, user and adaptation models in order to perform semantic retrieval adaptation. In Linckels, Sack and Meinel (2007), learning objects are adaptively retrieved exclusively using such inferences of Description Logics, which is possible due to the semantically marked up documents. In Tran, et al. (2008) it is even proposed to have one domain ontology that encompasses all the different aspects of Adaptive Hypermedia, such as the user or the task. Additionally, a model for the specification of adaptation rules upon this ontology is suggested to capture the adaptive behaviour in a declarative manner.

However, the use of such detailed concept, metadata and rule modelling requirements has confined many of the presented techniques to rather small-scale applications. More “lightweight” semantic techniques are proposed in Fernandez et al. (2008), where documents are semantically annotated using a combination of keyword frequencies, as well as semantically related documents’ keyword frequencies. Furthermore, their ranking algorithm makes use of a combination of conceptual ranking together with standard keyword ranking in order to lift the semantic techniques to a potentially large scale. While such semantic retrieval approaches do not have any explicit representation of user needs, there exists great potential to utilise or modify such techniques for adaptation and personalisation. An example of such personalised semantic techniques is shown by Cantador et al. (2008), where a semantic user model is created from the concept annotations that are associated with documents viewed by the user. Additionally, the user model concept weights are constantly updated depending on the frequency of user interactions with associated documents. The document ranking then makes use of these different concept weights in order to provide a conceptual ranking that compares the user preference vector to the document metadata vector. By combining a standard keyword search score with such a conceptual ranking, search results can be shown to be more personally relevant to querying users. However, the single axis of adaptation is the adaptation towards previously shown interests, resulting in an improved ranked list only.

4.3. Summary and Critique

Due to the inherent differences between statistical and metadata-based retrieval approaches, various techniques have been developed to add adaptivity and personalisation to the retrieval process. A summary of the surveyed techniques can be found in Table 2.

⁴ <http://www.jboss.org/drools>

⁵ <http://www.w3.org/TR/rdf-schema>

⁶ <http://www.w3.org/2004/OWL>

Table 2. Summary of Retrieval Adaptation Techniques

	Modified Statistical Ranking	Combined Statistical Ranking	Adaptive Hypermedia	Semantic Techniques
Document Index	Keyword Index	Keyword Index	Concept-Content mapping	Concept-Content mapping
Scale	Large Open Corpus	Large Open Corpus	Closed Corpus/ Small Open Corpus	Closed Corpus/ Small-Medium Open Corpus
Domain Model	ODP categories	ODP categories	Bespoke Domain Concept Model	Domain Concept Ontology
User Model	ODP categories, Personal Information Repository	ODP categories, Keyword History	Overlay Knowledge Model, User Preferences	Overlay Knowledge Model, User preferences
Metadata Requirements	Automatic topic categorisation	Automatic topic categorisation	Rich content metadata, Concept- content mapping	Rich content metadata, Concept- content mapping, Ontology-encoded knowledge bases
Adaptation algorithm	Modified HITS/ PageRank/ BM25	Reranking of initial statistical result, Category similarity calculation	Rule-based application of strategy/narrative	Semantic reasoning, Combined Statistical-Semantic Ranking

From a statistical IR point of view, it has been suggested to either *modify* traditional ranking algorithms or to *combine* original scores with an additional personalised score. Metadata-based approaches can be divided into techniques that stem from the fields of *Adaptive Hypermedia* and the *Semantic Web*.

In the case of *modified ranking algorithms*, traditional scoring formulas are either extended with additional parameters or they are biased towards more personally-relevant information. Personal relevance judgements, as well as directory structures are often employed to categorise both documents and users in order to calculate similarities across them. Additionally, personal desktop information is utilised to gather larger volumes of data that can be used to find personally relevant search results. Statistical keyword-based measures are used to find similarities between different sets of documents and user preferences. This constitutes a clear advantage of these types of systems, since the indexing of documents and their classification can be applied on a large scale.

However, the surveyed classifications occur on a very broad scale, as they often employ only a small subset of ODP categories. Additionally, the user model is confined to a simple set of (weighted) keywords, which is lacking additional semantics in order to infer more refined adaptation strategies. Especially in the case of an initial short keyword query, this type of technique does not take into account that additional documents beyond the provided set of terms are relevant to the user. There is a clear lack of a strategy or narrative for adapting the results to suit the particular information intent or need of the user.

Similar conclusions can be inferred for techniques that *combine* an original ranking function with additional personalisation features. These approaches often rerank an initial set of results according to previous user interactions or relevance judgements. ODP categories and users' query histories are used in order to create weighted user profiles. This fully automated process is again scalable for large open corpus collections, which constitutes the main advantage of these techniques.

However, due to the inherent reliance on an original, non-personalised set of results, it can be argued that a large set of personally relevant results is neglected in the initial retrieval. Even if the result list is reranked successfully, these systems are not diversifying the results according to users' actual information needs.

In the field of *Adaptive Hypermedia*, rich domain and user models are used in order to retrieve personally relevant information. Due to the fact that many systems have been developed for the field of e-learning, user models have often focussed on modelling a user's knowledge of the domain. Together with specific user preferences (e.g. learning styles), this knowledge model is used by systems to adaptively retrieve documents that a user should examine in order to fill a knowledge gap. The adaptation usually occurs on a conceptual level, with relationships between concepts being defined in the domain concept model. These systems allow a multi-dimensional adaptation that can take utilise multiple user attributes for personalisation.

However, since these systems operate on a concept level, a mapping has to be created between documents and concepts. This has restricted most AH applications to a closed corpus domain, with certain open-corpus solutions requiring considerable indexing effort. Furthermore, the retrieval of concepts is usually not initialised by a user query, as most educational systems focus on delivering a personalised course rather than a query response. This leads to the open issue of how to apply the developed adaptation techniques to the often sparse sets of user keyword queries.

Semantic techniques fully move towards semantic technologies such as ontologies for domain and user models. Additionally, adaptation rules can be encoded into ontologies using description logics and new query languages can be used to reason across semantic knowledge bases. Interoperability is the clear advantage of these techniques, as such systems might be able to share domain, user and adaptation models. Such techniques can be used for both adapting a conventional retrieval algorithm (i.e. perform PIR), as well as for reasoning about results and adaptively retrieving additional resources (i.e. perform AH).

Again due to the high construction costs for such systems (especially the concept-content indexing), closed corpus domains have dominated most of the research in this area. However, more "lightweight" semantic techniques that require less metadata and that include some statistical elements seem to improve this scalability issue, with the downside of focusing on user interests only.

A potentially promising variant of semantic techniques lies in the field of open hypermedia systems (OHSs). These are not specifically adaptation- or personalisation-oriented, but could be combined with such systems to provide the most lightweight personalisation strategy of all. Some OHSs of all apply a late-binding of link structure (Brailsford, 1999) (Davis et al., 1992) and, in some cases, of search results, relying on the definition of link structure computations which firstly establish the scope of the computation (and thus the eligibility of any given object to be the source of a link), and secondly to determine the destination of the link, frequently with a retrieval operation over an information system. If these link structure computations can be specified in different retrieval query languages (such as SQL), it becomes possible to create links which are essentially scoped searches, where the context of the search is known because the user invokes the search by clicking on the link anchor (Verbyla and Ashman, 1994). While the early OHSs did not by default include user-specific data in the destination query, it was certainly feasible to do so in the link structure computations, and required only that the author of the computations incorporate the user-specific data into either or both of the context computation and the destination computation. This ultimately lightweight approach requires no system-specific metadata, delivering results according to the capabilities of the system to which it sends computations. This in turn means that the capabilities of the open hypermedia system are governed by the capabilities of the third-party systems to which it sends its requests.

4.4. Comparison across Retrieval Adaptation techniques

In terms of *document indexing*, systems that use modified ranking algorithms or that combine original search results with a reranking module utilise standard keyword frequency measures. This can be calculated automatically and can therefore be applied to large document sets. For an Adaptive Hypermedia system, documents have to be mapped to domain concepts, which is often done manually. IR systems therefore clearly outperform Adaptive Hypermedia applications in terms of efficiency. However, current indexing techniques do not cover additional document characteristics, such as its suitability for different types of users and contexts (e.g. novice/expert, time constraints, etc.). These types of document indices can often

only be created manually or semi-automatically. Adaptive Hypermedia has its strength in handling these different dimensions of adaptivity that could be applied across a document base.

The *domain model* of statistical systems most often consist of ODP categories in order to classify and personalise ranking results. These techniques have the clear advantage of being applicable to large collections of documents without human interaction. However, the strength of Adaptive Hypermedia techniques lies in the more detailed specification of domain models using either bespoke concept models or general ontologies. Although these models are more labour intensive to create, they allow adaptation strategies to reason over which concepts should be included in the retrieval. Especially in combination with a semantic user model, these domain ontologies can infer concept dependencies that are not captured adequately by keyword similarity measures.

Similar conclusions can be drawn from the surveyed systems in terms of *user models*. Statistical techniques rely on the machine processing of either ODP category classification or personal information repository mining. Weighted keywords denote user preferences, with minor keyword relationships being established using co-occurrence. The strength again lies in the processing of large sets of documents, which can lead to the mining of potentially large amounts of personally-relevant keywords. However, a semantically richer user model allows systems to infer which additional relevant information should be presented to the user in order to address the current information need. Additionally, the surveyed AH systems explore a multitude of dimensions of adaptation, using an increased number of user variables (e.g. prior knowledge, learning styles, time available). By contrast, the sole dimension explored by statistical techniques is based upon user interests that have been mined from previous search interactions.

As a consequence, the *metadata requirements* for statistical IR techniques can be computed efficiently, as techniques such as information extraction and topic categorisation can be fully automated. The disadvantage of Adaptive Hypermedia (including semantic-based techniques) lies in the reliance of not only marked up documents, but since the reasoning occurs on a concept level, a concept-to-content mapping needs to be created. Since current systems require a rigid annotation and indexing of these resources, manual effort is usually required to produce the desired quality of metadata. A better solution would loosen these requirements in order to provide more lightweight reasoning solutions. Additionally, this would allow the integration of (semi-)automatic IR indexing techniques in order to process larger open corpora.

The *adaptation algorithms* used by the statistical systems either integrate personalised features directly into the ranking function or they combine an initial score with a separate personalisation score. In both cases, they are applicable on large document bases, relying mostly on either term frequency measures or category similarity calculations. As mentioned before, combined score techniques only rerank an initial set of results that might be missing valuable information. It is therefore advisable to include personalisation features directly in the ranking formula in order to find personally relevant documents across the whole collection. Adaptive Hypermedia techniques make use of rule-based strategies or narratives that capture particular adaptation techniques. Using these rules, it is possible to extend the retrieval beyond an initial set of user defined keywords by examining which additional resources address a user's information need. An improved solution could apply these techniques to an open corpus space by combining the strategy model (or semantic reasoning in the case of semantic techniques) with loose concept indexing of large document corpora.

5. Adaptive Composition & Presentation

The set of results returned after the adaptive retrieval stage can be personalised further during the composition and presentation stage. Due to the variety of algorithms and methods presented in the previous section, the type of output results differ significantly from system to system. Therefore, an array of techniques has been developed to perform different types of adaptive result composition and presentation. Again, these methods can be divided into statistical and metadata-based techniques, depending on the algorithm that generated the results.

5.1. Statistical techniques

Traditionally, ranked lists have been the preferred method of displaying statistical information retrieval results. This is reinforced by the fact that most research in the IR field, including result adaptation and

personalisation, is concerned with improving single-valued relevance scores, which can only be displayed in the ranked list format. Although this score might be calculated using several features/algorithms (including personalisation features), a single aggregated score is generally calculated in order to simply compare the different values for ranking. However, recent research has attempted to provide alternative composition and presentation techniques in order to provide more personalised, adaptive and diversified results. A selection of such systems is surveyed below.

First of all, the composition (merging) of search results from a set of multiple search queries has been explored by several researchers. The idea behind such techniques lies in the broadening of search results in order to either focus a user's search towards more precise information needs or to improve the performance of personalised reranking techniques. For example, in the metasearch engine Inquirus2 overviewed earlier, Glover et al. (2001) use the results retrieved by various search engines and compose these into a single ranked list. As opposed to typical metasearch engines, they not only consider the titles, summaries and URLs for the rank merging, but the whole pages returned by each engine. The ranking is based on multi-attribute utility theory, which takes into account several factors, depending on which need category was chosen by the user. The different preferences are used in an additive value function, which combines the different metadata fields that are available. For example, the indicated preference for "current events" would put a 60% emphasis on "TopicalRelevance" and 40% on "DaysOld". Similarly, Radlinski and Dumais (2006) also propose the diversification of search results through the merging of multiple result sets. However, they propose to generate the set of multiple queries by determining related queries from a large sample of query logs. In combination with reranking techniques proposed in Teevan, et al. (2005), the diversified results are shown to provide improved personalised rankings.

The idea of retrieving diversified search results is also proposed in the meta-search systems in Sushmita, Joho and Lalmas (2010) and Thomas et al. (2010). However, these systems make use of an "aggregated search interface" in order to compose and visually present a more diverse set of results. Similar to the diversification techniques above, an initial query is sent to several information sources in order to retrieve diversified search results. The resulting documents are then not simply combined into a single ranked list, but they are displayed in a separate panel for each information source on the same "aggregated" results page (similar to Yahoo alpha⁷). Experimental results show that these prototype systems enable users to look at more diverse results, select more items to complete their tasks, and that they are generally perceived to be superior to a standard ranked list system. The idea of such systems hence lies in the immediate visualisation of more diversified search results, as opposed to just attempting the improvement of a single merged result ranking.

Another approach of using statistical methods to adaptively compose and display results is the concept of clustering. In Xu, Jin and Lau (2009), a user query is sent to a third-party web search engine to retrieve *N* number of results. This set is then clustered into different topics using standard document clustering techniques. Following this step, the main topics are allocated a display panel in the visual interface. The size and location of each panel depends on the size and importance of the search results contained in each cluster. A user can then either click directly on one of the search results from a cluster, or expand a particular cluster to display the full results of the chosen topic. Yippy⁸ takes a similar approach by categorising search results into folders and subfolders. A user can then expand a certain folder/subfolder to refine the search, therefore making the interface adaptive to user interactions. In Truran, Goulding and Ashman (2005), multiple clicks ("co-selections") on a set of search results are interpreted as indicating mutual relevance. By mining such relationships, their system is able to aggregate search results of ambiguous queries into a set of clusters that can help users sort through the different query senses. Although this usage of collective intelligence has so far focussed on non-personalised aggregations, it is worth considering the creation of class/cohort specific systems using this technique.

Some research has been conducted in improving the traditional result presentations by providing the user with increased information about the retrieved resources. In Psarras and Jose (2006), an improved summarisation system is proposed, which performs adaptive query-biased summarisation. These summaries are presented with the traditional ranked list and are shown to improve users' relevance judgement during result browsing. The system is implemented as a recommendation portal, which

⁷ <http://au.alpha.yahoo.com/>

⁸ <http://www.yippy.com>

adaptively presents relevant documents to a user based on previous searches and result visits. In White, Jose and Ruthven (2003), the “WebDocSum” interface similarly provides users with an improved summary in the form of a summary window. When a user moves the mouse over one of the query results, this window displays a summary for the document. In addition to the usual fields such as title and summary sentences, it provides the user with the number of outlinks on the page, the first non-text object and the document size. Similarly, in Joho and Jose (2008) an evaluation of 4 different search result presentations is performed. A baseline ranked list is compared to (i) a system that presents top ranking sentences along with each result, (ii) a system that shows a thumbnail image of each result document (i.e. a screenshot of the actual document) and (iii) a system that presents both top ranking sentences and screenshots with each result. Although neither of the two research works above describe adaptive result presentation, it is argued that differences could be noticed among users with different search experience. More specifically, it is noted that less experienced users might benefit from the systems that provide extra information for the search results in order to make better relevance judgements. Furthermore, the search interface should be made adaptive to the particular task, context and user experience to offer the right and appropriate assistance at any given time. In related research, Villa et al. (2009) provide an adaptive “aspectual” search presentation that allows users to model search subtasks. For each aspect, users can have a separate panel with its own history, undo history, current query, search results, etc. This adaptive presentation allows for the completion of complex information needs, which require users to search for multiple aspects within the overall task.

While the techniques and approaches above attempt to better organise, compose and present an initial set of search results, Bhavnani et al. (2003) propose to use domain knowledge from experts in order to develop actual search strategies that can help novice searchers find information more effectively and efficiently. They developed the idea of Strategy Hubs, which provide initial selection categories, which in their case are related to medical conditions. A user can initially choose from a selection of diseases, followed by a selection of subcategories, such as “Treatment” or “Diagnosis”. In a second step, the hub provides specific search strategies about how the user should find information related to certain topics. Additionally, for each of the strategy steps, links are provided for reliable sources that are known to provide good information. The strategies consist of a series of sequenced steps that have been identified by experts. In terms of presentation, a dual-frame design has been chosen, which displays the different steps in the upper-frame and the actual content pages in the lower frame. It is argued that this design provides a consistent user interface, supporting novice users in their perception of the overall strategy. Since novices have greater difficulties in identifying sub-goals when searching for comprehensive answers, the strategy hub can support users by guiding them towards a more structured way of searching.

5.2. Metadata-based techniques

In the presence of rich metadata, faceted rankings have become a popular way of composing and presenting IR results in a more easily comprehensible manner. Such systems allow users to search for information through the specification of more refined attributes than just simple keywords (Yee, et al., 2003). For example, in Figure 6 a user has refined the image search according to “Location: Asia” and “Shapes, Colors, and Materials: fabrics”, resulting in a narrowed down set of image results. Such metadata attributes are typically added manually to individual items in a collection, although they can be extracted automatically to a certain extent. Such a faceted way of ranking search results is also proposed in Teevan, et al. (2008), where facets are described to “represent a dimension that can be used to organise information”. It is argued that an adaptive process could choose between different facets depending on the user task and context, as well as the document domain. By selecting multiple rank facets (using different document attributes), a user could adapt the result presentation towards a more personalised view of relevancy scores. This idea is also proposed in Tvarozek and Bielikova (2007), where only selected metadata fields are used to show/rank the most relevant attributes of the search results. Furthermore, the authors propose an adaptive version of their faceted browser, which provides automatic facet selection based on user preferences, global attribute relevance and inter-attribute relationships. Additionally, facets can be adaptively ordered (based on their estimated relevance), annotated (e.g. using tooltips), and recommended based on particular restrictions (e.g. IT companies being recommended to an IT consultant). Similarly, Zhang and Zhang (2010) propose to recommend document facet-value pairs to users and to incorporate the selected values into the retrieval models. Experimental results show that for a corpus of semi-structured text documents, a non-boolean retrieval model performs more effectively.



Figure 6. Example of a faceted search interface (Yee, et al., 2003)

In contrast to traditional IR systems, the field of Adaptive Hypermedia (AH) has been inherently focused on providing users with adaptive result compositions and presentations. Due to the availability of rich user and domain models, coupled with bespoke content metadata, ideas in this field focus around composing appropriate information sequences and navigations and displaying these using personalised presentations.

As mentioned earlier, these systems have generally focused on delivering educational material to students in order to provide adaptive e-learning courses. For example, as mentioned in the adaptive retrieval section, the KBS Hyperbook system described in Henze and Nejd1 (2001) performs concept-level adaptation by making use of domain and user models. The selected concepts are sequenced according to prerequisite requirements, creating an order of links that guides a student towards the next best document to view. It is not left to the student to sort through a ranked list of documents, as the system advises an appropriate path through the document space through link sorting/hiding. Since the actual documents have been indexed with concepts from the domain model, any open corpus document can be sequenced appropriately. Similarly, the idea of providing appropriate sequences of documents is shown in Farrell, Liburd and Thomas (2004), where XML search results are mapped to topics and then sequenced according to particular concept domain rules (e.g. to teach more basic information first). Additionally, the actual objects within topics are sequenced according to an “Instructional Role Sequence” (for example introductions are sequenced before concept procedures and conclusions). However, the sequencing service is embedded into the adaptive system rather than implemented as a separate adaptation model. Also, no personalisation is provided apart from the possibility to perform query refinement.

The generation of personalised learning sequences is taken further in the APeLS system (Conlan et al., 2002) (Conlan and Wade, 2004), where the separation between the core adaptive engine and the sequencing service is proposed in the form of a narrative model. This narrative reflects a didactical ordering that can be specifically adapted to the current user task, context and preferences. By applying this adaptive narrative, the system can provide a personalised result sequence based on the particular user knowledge and preference levels, hence guiding the user through the document space (again through link generation, link sorting, etc.). Although the system has been initially conceived to work over a closed corpus, it would be possible to index and consequently integrate open corpus documents, as in the case for the KBS Hyperbook.

The ideas of composing and sequencing information have also been explored using a combination of ontologies and description logics. For example, Karam, et al. (2007) assemble learning objects by inferring the best “composition flow” using the current user knowledge state and the domain ontology. Description logic is used to solve the “concept-covering problem”, which corresponds to the knowledge need of the

current user. Similarly, in Geurts, et al. (2003) ontologies are used in order to provide a "structured progression" through retrieved results. The concept of narrative units is proposed, which are used to construct the complete result presentation in a structured order. Each retrieved semantic unit has rules associated with it, which dictate the information that should follow the current unit. As a consequence, after the application of the complete set of rules, the result is a structured progression through a semantic graph.

In addition to composing such adaptive navigations, AH systems often use adaptive presentation techniques such as link adaptation in order to provide users with presentational hints. For example, by colouring or hiding selected links, Smits and De Bra (2011) provide hints to users about the relative suitability of particular knowledge items. With growing knowledge of a user, different links either become available or get coloured to symbolise their "readiness" to the user. The knowledge modelling in such applications relies on rich domain models, which can be used as a basis for overlay user models in order to infer the suitability of particular content. In Hsiao, et al., (2009) and Hsiao, et al., (2010), a user's progress through a course-test system is tracked in order to provide adaptive link annotations. Such annotations provide students with hints about which task to try next and also visualise how often a certain quiz has been attempted already. In Figure 7, the target icon in the menu (left) presents the growth of student knowledge (shown by the number arrows) and the relevance of the topic to the current course goal (shown by the colour intensity of the target, ranging from faded to strong intensity). Such techniques are shown to lead users to attempt more course tests and to have higher success rates, as adaptive link annotations seem to have a high motivational effect in such applications.

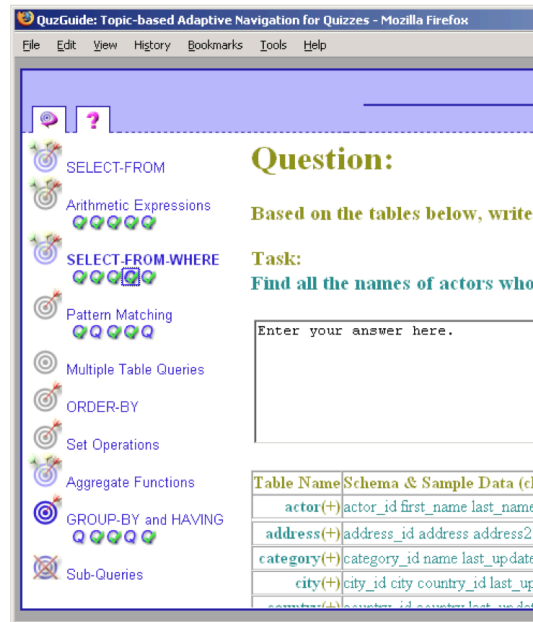


Figure 7. Example of link annotation (Hsiao et. al., 2009)

In Jovanović, et al. (2006), sets of knowledge items are grouped into an annotated tree of links and link annotations are provided to show users which documents are most appropriate based on prerequisite requirements. Upon selecting one of the links, the system then generates a new assembly of learning content based on the selected topic.

By contrast, rather than showing a user a full sequence of documents, the systems in Henze (2005) as well as in Smith-Atakan and Blandford (2003) consist of an adaptive presentation in the form of related information that are apart from the currently viewed content. The former presents the current best pages on the left of the current window, whereas the latter system, called ML Tutor, provides these in a separate window. In Brusilovsky, et al. (2004), The Knowledge Sea II system provides a user with a map where similar documents are placed in adjacent cells. It provides social navigation by using visual cues based on an individual user's browsing history combined with all other system users. The popularity of a particular document is highlighted both on an individual and on an overall level, guiding a user towards popular

documents that he/she has yet to visit. Additionally, users can provide annotations, such as positive or negative feedback, which can further help fellow users find interesting documents. However, as noted by the authors, although the system is very efficient in adding open corpus documents, it is lacking a strong navigation support, especially for an educational system.

Further adaptive presentation techniques include scaling, where important information is highlighted through increasing the size of relevant content, or stretch text, whereby less relevant content is only represented by placeholders (Tsandilas and schraefel, 2004).

5.3. Summary and Critique

The adaptive composition and presentation of information has been studied by both fields of *Personalised Information Retrieval* and *Adaptive Hypermedia*. Each approach is trying to overcome the information overload problem by grouping, sequencing and presenting documents in a coherent manner. A summary of the different techniques and approaches is provided in Table 3.

Table 3. Summary of Adaptive Composition and Presentation Techniques

	Adaptive Composition		Adaptive Presentation	
	Information Retrieval	Adaptive Hypermedia	Information Retrieval	Adaptive Hypermedia
Adaptive Behaviour	Faceted Ranking, Search Aggregation, Result Clustering	Adaptive navigation, Concept/content sequencing, Link ordering, Map-based indexing	Adaptive result summarisation, Adaptive cluster presentation	Link annotation, Link colouring, Map-based presentation, Stretch text, Scaling
User Involvement	User intent elicitation, Facet preference selection	none	Cluster selection Aspect creation	none
User Model	none	User Knowledge/ Preferences	Result summarisation preferences	User Knowledge/ Preferences
Metadata Requirements	Intention-facet relationship model	Concept-content index, Keyword similarity	none	Concept-content index, Keyword similarity
Adaptation algorithm	Statistical document/ keyword analysis	Rule-based (ECA), Social-based, Semantic reasoning	Statistical document/ keyword analysis	Rule-based (ECA), Self-organised social mapping
Scale	Large open corpus	Closed Corpus/Small-Medium open corpus	Large open corpus	Closed Corpus/Small-Medium open corpus

In the field of *Personalised Information Retrieval*, most research has generally focussed on adapting the result compositions and presentations by reranking initial search results. The presented techniques move away only slightly from the current ranked list paradigm by grouping results into clusters or by ranking the results according to particular facets. Additionally, the current features of ranked lists (e.g. result snippets) have been adapted in order to provide more personalised result presentations. As in the previous sections, statistical document analysis techniques provide this improved visualisation. Document clustering and snippet/summary improvements can be applied on a large scale, making these techniques attractive to be applied in current web search.

The concept of result diversification and aggregated search moves away slightly from the single query and single ranked list paradigms in order to present users with a greater breadth of search results. However, neither technique makes use of adaptive strategies in order to choose the right type of query diversification and information source for a particular user, task or context.

Additionally, the notion of information sequencing has yet to be addressed in order to assist the information searcher more adequately. Search strategies or procedures are not supported, leaving users having to filter through large ranked lists in order to satisfy their information need. Also, the presented systems do not take into account that users might have previously acquired particular information about a certain subject, which would ideally decrease the future relevancy of documents that cover this part of the knowledge space.

On the other hand, due to their inherently adaptive behaviour, *Adaptive Hypermedia* systems address this user guidance using various approaches. Concepts and content are composed into sequences of coherent (learning) paths, which assist the user in finding additional documents for their initial information need. By making use of user models that contain current preferences and knowledge levels, this path can be adapted to form a personalised response for a particular user. Such systems also have a strong focus on applying adaptive presentation techniques, effectively guiding users through content using adaptive hints.

However, due to the high reliance on metadata and concept-content indexing, most systems have been confined to small closed corpus spaces. The only system which does not require careful metadata indexing is the Knowledge Sea II system by Brusilovsky et al (2004). However, it is argued that the social navigation component in this system does not provide the same quality in terms of user guidance (Brusilovsky and Henze 2007). It is therefore important to carefully balance the issues of scalability and navigation support in order to provide an open corpus system with adequate levels of user guidance and personalisation.

5.4. Comparison across Adaptive Composition & Presentation techniques

The survey of the different systems reveals a variety of techniques to add *adaptive behaviour* to the composition and presentation stage. In the field of Information Retrieval, the focus has been to slightly evolve the current ranked list paradigm by adding either cluster visualisation, faceted ranking or improved snippet/summary generation. Although these techniques are easily applicable on a large scale, they do not provide any of the advanced adaptation techniques that are present in Adaptive Hypermedia systems. Current AH approaches address the very important aspect of concept/content sequencing, which provides more guidance for querying users. Furthermore, the presentation adaptation of AH systems often provides additional visual cues (such as link colouring or hiding) to the user in order to show the personal suitability of particular documents. This adaptive behaviour represents one of the true strengths of AH systems and could therefore provide excellent extension possibilities to current IR systems in order to make use of their bulk-processing in a more personalised manner.

In terms of personalisation, IR and AH systems have both increasingly investigated the concept of facet preferences in order to provide more suitable rankings for a particular user. *User involvement* is still the most valuable input for IR systems in order to infer precise and context-sensitive relevance scores. AH techniques make use of *user models* to store a user's prior knowledge and additional preferences in order to predict appropriate information relevance. Although this addition would prove very useful in IR systems as well, it constitutes a great challenge to mine such information using standard data mining techniques. In order to retrieve this prior knowledge information from open web data, one would first need to map users' page visits to a domain model and consequently estimate a user's information gain about the appropriate concept. This problem represents a fundamental issue in implicit feedback techniques, as a user's browsing activity does not necessarily correspond to learning about the documents' topics. Different approaches can be applied to address this problem. For example, one possibility is to assign suitably low value-adds for page visits in order to avoid overestimating a user's experience. Alternatively, rather than looking at a user's search or browsing sessions (page counts, word counts), another approach could employ public user profile mining. Once a user has been identified, user profile data from social networking sites such as Facebook or Twitter could be mined and consequently utilised within the personalised application (Abel et al., 2011).

The *metadata requirements* for the surveyed IR systems are comparably low, as the main focus has been on automatic document classification and result summarisation. This independence from markup data makes

these systems very suitable for large scale corpora such as the web. However, the emerging concept of faceted ranking will require considerable amounts of metadata to order documents according to particular preferences. From an AH perspective, this dependency on sufficient metadata has been apparent since the earliest systems. Documents need to be linked to domain concepts in order to reason about the suitability for a particular user. However, systems such as the Knowledge Sea system by Brusilovsky et al. (2004) have made initial steps towards a more open AH system. The compromise between reliance on metadata, user guidance and scalability has to be chosen carefully in order to provide the right level of adaptivity for the particular task and context.

As mentioned in the previous sections, the *adaptation algorithms* used by IR systems have built on statistical keyword similarity and clustering measures in order to provide result compositions. User relevance feedback is captured and added to the ranking formula in order to update the rankings/clusters. AH systems on the other hand have mainly made use of rule-based algorithms to compose suitable result flows. Ontologies are being used increasingly, which leads to the encoding of adaptation rules in formal languages such as description logics. Since these rules are currently not being created automatically, a scalability issue arises again for AH systems. A solution to this problem would be to use high-level adaptation strategies that could be complemented by automatic IR processing capabilities.

6. Conclusions

The fields of Personalised Information Retrieval (PIR) and Adaptive Hypermedia (AH) have each recognised the challenge of adaptive and personalised information delivery. However, due to their inherent conceptual differences (presented in section 2), the techniques and technologies have varied substantially between the two approaches (described in sections 3, 4, 5).

Section 6.1 first presents several conclusions that can be drawn in terms of the user dimensions (i.e. characteristics) addressed by PIR and AH systems. Section 6.2 draws a number of conclusions about the different adaptation techniques that PIR and AH have applied to adapt to such dimensions. In addition, while the techniques of both PIR and AH have progressed significantly, there are several remaining personalisation challenges. Section 6.3 presents a number of such personalisation challenges that current PIR and AH systems may be facing. Lastly, in order to overcome some of the identified weaknesses, section 6.4 investigates the potential benefits of combining PIR and AH techniques in a hybridised approach.

6.1. User dimensions

In the field of PIR, adaptation and personalisation techniques have predominantly focused on the statistical analysis of historical usage and corpus patterns, using for example past queries, query refinements or user clicks (see section 3.1 and 4.3). The analysis of such usage patterns has typically mined sets of user interests, which can then be used for narrowing future retrievals towards related information.

As discussed in section 3.1 and 4.3, the major advantage of such techniques has been shown to lie in their scalability, as the algorithms are mainly focussed on automatic processing of large volumes of data. Search logs, personal information repositories, as well as directory structures such as the ODP have been used increasingly in order to personalise search results by categorising both users and documents. Some systems have even mined minor semantic relationships between documents, concepts and users by analysing query histories and consequent click behaviours, including e.g. co-selections on a set of search results such as in Truran et al. (2005).

However, it is unavoidable that such techniques may introduce substantial noise, especially due to the variety of user information contexts and (sometimes extemporary) needs. One way to overcome such noise (e.g. misrepresentation of user interests) is through scrutability and user control, whereby users are able to manually review and adjust the current user model. Moreover, it is crucial to recognise a user's current task in order to provide not only personalised information, but also contextualised personalisation. Initial attempts have been made in order to tackle this issue, for example through contextualising a user need by using only sets of consecutive queries for personalisation (Xiang et al., 2010). This ensures that for example different usage patterns (e.g. casual browsing) are not affected by irrelevant long-term user needs. Similarly, Cantador et al. (2008) model short-term "context models" in a similar fashion to the more long-

term user models (i.e. by mapping keyword vectors to semantic concepts), but using a strong decay factor to fade out older concepts. This context model is then compared to the user model in order to select the relevant subset of a user's long-term interests for personalisation. This ensures that a user's long-term interests are taken into account, while not over personalising results based on contextually irrelevant information.

The fact that PIR techniques typically base personalised relevance estimation solely on previous user interests also constitutes a very narrow focus of adaptivity articulation for such systems. Since users constantly interact with systems in order to fill particular knowledge gaps, it is crucial to consider that current information needs depend on a number of characteristics, such as a user's task or current knowledge state.

AH techniques have inherently focussed on capturing and using such additional dimensions of adaptation (e.g. prior knowledge, cognitive/learning style) in order to provide user assistance and guidance (see section 4.3 and 5.3). Such multi-model techniques have mostly focussed on instrumenting the usage during a domain-specific user session and are therefore less ambivalent or susceptible to noise. Moreover, as AH techniques are inherently based around structured conceptual models rather than unstructured keyword histories, the aforementioned scrutability and user control can be achieved more easily through the development of model manipulation interfaces (Bakalov et al, 2010).

However, the rich conceptual modelling techniques of AH systems have typically been confined to narrow domains such as e-learning and it remains an open challenge to broaden some of the techniques to larger content bases. In particular, it may be necessary to also use statistical usage methods (similar to the presented PIR techniques) in order to track such information across open-corpus domains. Initial attempts have been made in the context of studying open-corpus novelty detection, whereby a user's domain knowledge is calculated using a knowledge accumulation methods based on previously viewed documents (Lin and Brusilovsky, 2011). The suitability of a new document can therefore not only be characterised by its similarity to user interests, but also by the estimated suitability to a user's current knowledge state. However, these novel techniques for opening up traditional AH knowledge modelling still need to be evaluated thoroughly in terms of their accuracy and usefulness in applied scenarios.

6.2. Adaptation techniques

In the field of PIR, adaptation techniques have most prominently focused on query expansion and result reranking. Such techniques have typically been based on statistical similarity measures using keyword-based user and document models (see sections 3.1 and 4.3). An array of different techniques has been proposed to expand, modify or trim initial user queries, as well as to bias retrieval algorithms towards statistical user models. The techniques have been shown to successfully retrieve focused result lists according to prior user interests and they have generally maintained the scalability of standard search systems.

However, the notion of adaptively choosing different query/retrieval adaptation techniques has been less explored to date. One area where such ideas have been actively researched is in the related field of Question Answering, where systems may involve query expansion based on synonyms, external thesauri or parsing a user's question with grammars of varying sophistication. The goal of such adaptive selections of multiple query adaptation techniques typically lies in broadening the potential pool of answers before proceeding to the actual answer extraction step (Hirschman and Gaizauskas, 2001). However, while such techniques could also be employed to broaden and diversify initial user queries, PIR systems typically revert to the presentation of their results using the conventional ranked list paradigm (see section 5.1).

This constitutes one of the main drawbacks of PIR approaches, as users are typically left having to filter through simple ranked lists (or possibly clusters of ranked lists) of potentially relevant/non-relevant documents. No explicit user guidance according to a strategy or narrative is provided across documents, a fact that is reinforced by the common IR batch-evaluation techniques that do not involve real-world users. This lack of narrative is one of the key distinctions between PIR and AH, and aligns with the distinction noted previously between searching (PIR) and browsing (AH) (see section 2). While PIR is almost entirely search behaviour, AH is characterised by browsing behaviour guided through a narrative defined previously by an expert.

By contrast, AH techniques have inherently focussed on providing such user guidance through external models to enable particular information seeking strategies. Adaptive result composition and sequencing are used to provide a flow of currently suitable information in order to provide “the right information at the right time” (see section 4.2). Moreover, as systems have typically used implicit queries (e.g. the statement of an intent/learning goal), AH techniques have focused on scoping the adaptive experience rather than adapting an initial keyword query. Many systems have been conceived for the application domain of e-learning, where user guidance can be provided by a domain expert (often a teacher/lecturer) through the encoding of a domain model. Adaptation strategies can then be applied on this model by defining rules upon relationships, such as prerequisite requirements. These rules can be applied in order to provide additional dimensions of adaptivity, such as localisation or context/task-based personalisation.

Moreover, AH systems also make use of a multitude of adaptive presentation techniques such as link colouring or annotation (see section 5.2), thereby providing a much more guided browsing experience. These presentation techniques present one of the true strengths of AH systems, as they can provide personalised hints without hiding information from users. This also highlights again the focus of AH systems on user navigation and interaction compared to the batch computation of ranked result lists.

However, due to the inherent reliance on refined concept indexing, most research has still been confined to very narrow domains, such as educational systems or cultural heritage libraries. Due to the closed nature such systems, most research in recent years has focussed on how to move towards open-corpus domains. Lightweight solutions to concept mining, indexing, reasoning and adaptation are required, which make use of both the bulk-processing capabilities of (P)IR and the adaptation and personalisation approaches of AH. The Semantic Web field has also introduced such increasingly scalable solutions, although they have mostly had the downside of using their semantic capabilities for the purpose of only improving ranked lists using the sole dimension of user interests (see section 4.2). Additionally, the Linked Data initiative⁹ promises a large-scale availability of structured data that could be used by AH and semantic systems. To date, over 32 billion triplets have been published already, which could help AH to overcome some of their scalability issues. However, it remains to be seen how effectively these datasets can be used by adaptive systems.

6.3. Personalisation Challenges

PIR and AH both face challenges in terms of providing personalised user experiences while still maintaining balanced, fair and ethical services. In particular, while the primary purpose of personalised systems is to find and present the most personally relevant information, systems have to ensure that users do not suffer from a loss of serendipity. This phenomenon, sometimes also referred to as “Tunnel Vision”, means that users may be presented with an information focus that is too narrow (Billsus and Pazzani, 2007).

PIR in particular could suffer from this problem, as the ranked list paradigm typically entails users not looking beyond the first page of results, as is already manifested in 'normal' search results as 'trust bias' (Joachims et al, 2005). This aspect could potentially lead to users failing completely to receive contrasting opinions outside of their own set of (previously exhibited) interests and preferences (recently referred to as “the filter bubble” (Pariser, 2011)). PIR systems therefore need to ensure that users are aware of the personalisation which has occurred (as well as the reasons why) and that such effects can be turned off if desired. A potential solution is to diversify the results on the first page in order to provide an overall picture of the information space (e.g. as presented by Radlinski and Dumais, (2006)).

The filtering problem is also particularly important in the context of possible censorship, as retrieval engines could deliberately bias search results for commercial or political reasons, and indeed such accusations have already been made (BBC, 2010). There is even a claim that "Personalization isn't really about search, but fine-tuning the relationship between users and advertisers" (Sherman, 2011). If users are unaware of this deliberate adaptation, they could be influenced subconsciously and therefore fail to get a balanced view of factual information. As opposed to more obvious types of information censorship (e.g. suspension of a website/service), the alteration of search results presents a much more covert (but similarly effective) method of information control.

⁹ <http://linkeddata.org/>

Similarly, AH systems need to make sure that users are aware of the personalised experience within the overall information set. As noted by Ashman et al. (2009), providing personalised presentations may have negative impacts on users' overall knowledge acquisition, as they could fail to independently create their own mindmap of the underlying information space. Users might simply follow the adaptive guidance without making conscious decisions about which information to view next. AH systems therefore need to maintain a fine balance between providing adaptive guidance as well as encouraging freedom to explore the information space independently.

Lastly, privacy, data confidentiality and data ownership are all major issues that are equally relevant to both PIR and AH systems. As mentioned previously, systems need to make sure that users are aware of the adaptation that has occurred, as well as which personal information the system has used to perform this adaptation. Both PIR and AH systems need to ensure that the information is held securely and that third-parties are unable to extract user-sensitive data. A good example of privacy-protected personalised search is presented in Chen, et al. (2011), where it is proposed to hold the complete user profile on the client side. Moreover, Chen, et al. present customisation possibilities that allow users to specify "forbidden nodes", which are user model topics (i.e. ODP classes) that should not be disclosed to the remote search engine. Only this pruned user model is then sent along with the original user query in order to generate personalised results.

6.4. Overall findings

Although both PIR and AH have attempted to address the same challenge of *delivering* personally relevant information, each approach has presented different strengths and weaknesses in terms of user dimensions and adaptation techniques.

First of all, PIR systems have been shown to typically only provide adaptation according to the narrow dimension of user interests. Moreover, the adaptation techniques have typically been confined to the simple alteration of search result rankings, which can lead to the previously described problems of over personalisation and low information diversity. In order to increase the breadth of adaptation capabilities of such systems, it is therefore perhaps crucial for PIR to embrace the notion of multi-dimensional adaptation that current AH systems provide. Moreover, the AH notion of adaptive guidance in terms of result composition and presentation could be beneficial to PIR systems in order to overcome the low user commitment in current search systems. By utilising adaptive composition and presentation techniques, PIR systems could potentially engage users into personalised search sessions and motivate them to subscribe to the notion of search as an interactive process.

Several weaknesses have also been shown for AH systems, most notably in terms of their strong reliance on rich metadata models for retrieving information. This characteristic of AH techniques has typically confined such systems to small-scale application domains such as e-learning. In order to overcome these weaknesses, AH systems need to embrace the power of statistical document analysis techniques that have been shown to successfully drive the adaptation in current PIR systems. Techniques such as keyword query expansion and selective information source selection could enable AH systems to provide their multidimensional adaptation across larger open-corpus domains.

However, it is only possible to provide these combined functionalities if the complete retrieval process is enhanced. Most research so far has focussed on providing adaptivity only during either the query adaptation stage, the retrieval stage or the composition/presentation stage. Very little attention has been devoted to providing a unified adaptation approach, which could encompass all aspects of the task of information retrieval and delivery. However, it is crucial to align the different stages into a coherent workflow in order to enable personalised guidance during the information composition stage.

The development of "hybridised" systems could potentially combine different techniques and technologies in order to provide such a unified adaptation. Section 7 presents an initial exploration of the combination of PIR and AH techniques and elaborates more clearly the possible affordances that such a hybridised approach might provide.

7. Hybridisation

As outlined in section 6.4, one way to combine a number of personalisation techniques into a coherent adaptation process could be through the development of a hybrid PIR-AH system. The requirements for such a system would lie in the appropriate selection and application of AH and PIR components. The kind of affordances that such adaptation components could provide can reach across i) query adaptation, ii) retrieval and iii) result composition and presentation. Firstly, section 7.1 explores how each of these stages can be influenced by both PIR and AH techniques. Secondly, section 7.2 presents an example of such a hybridised approach in a real-world application scenario and explores the potential benefits of combining the complementary affordances of PIR and AH.

7.1. Potential combinations of PIR and AH techniques

First of all, multiple models can contribute towards the adaptation of the user's query, including both AH-type metadata models about the user's preferences and context, as well as PIR-type models of search histories. Moreover, AH strategies can be applied in order to choose between different types of query adaptation (e.g. query expansion, trimming, etc.) depending on the different model states. This can also include the generation of a set of multiple queries of varying detail and complexity in order to maximise the diversity of results. In particular, this diversification can aid the later composition and presentation states by providing a broader range of information related to the current topic of interest. Additionally, it is important to constantly update and reassess the various model states by monitoring user behaviour such as page views, dwell time or explicit relevance feedback.

In addition to query adaptation, the retrieval of information can be adapted using statistical PIR techniques in order to bias results towards personally relevant information. Such personalised results could also be mapped to AH models using statistical classification. This again allows the final composition and presentation stage to better guide users across the various results.

Finally, AH components can generate the navigation across the retrieved content based on the current user model states. Again, this needs to be done iteratively because user needs, preferences and knowledge are changing constantly during the information exploration. The models therefore need to be reassessed constantly using the knowledge mined from user feedback.

In conclusion, PIR components can assist during both the query adaptation and the personalised content selection. Similarly, AH needs to add in the ways in which it chooses the different query adaptation techniques (e.g. expanding multiple queries in a particular order), as well as in terms of the way it generates the adaptive navigation. Additionally, AH needs to provide immediate feedback to the various models in order to affect and adjust the overall adaptation towards the current context.

However, there remain many challenges towards achieving such integrated adaptation and personalisation. In particular, it is of paramount importance that the various adaptation stages are coordinated in terms of end-to-end effectiveness. If such a harmonized combination of techniques is not taken into consideration, the various adaptation effects could potentially neutralise each other or in some cases even be detrimental towards the overall system performance. For example, it might be desirable that an application does not perform adaptation on the same characteristics twice, as this might skew or over blow the results too much. This could be the case if a system personalises a query based on a particular characteristic and then also performs personalised retrieval based on the exact same attributes. Similarly, if a system has a broad-type strategy across the result navigation (e.g. to give a user as much choice as possible), it would not be advisable to perform too much focused personalisation during the earlier process stages.

The key principle towards the successful application of the proposed hybridisation hence lies in the joined-up thinking between the various adaptation characteristics and the understanding and managing of the trade-offs between techniques. Rather than arbitrarily combining multiple adaptation capabilities, system designers need to develop an overarching strategy that takes into account the system's application context and goals. The various adaptation techniques need to be coordinated according to this overall strategy in order to maximise the complementary affordances.

7.2. Application scenario

There are a number of areas where such hybrid systems could be applied in real-world scenarios, for example as information portals for companies and organisations. Such an information portal could be used in the area of customer support, which will be discussed in this section. Customer support systems aim at providing users with a variety of product information through functionality descriptions, guided tutorials or troubleshooting support. Such systems typically provide one-size-fits-all approaches, forcing users to either browse through long product manuals, large frequently asked question sections or even third-party community forums. Search systems across these resources often consist of non-adaptive IR engines that return the same ranked list for each user with the same keyword query. If a user's search for information using these resources is unsuccessful, a costly customer support call needs to be handled by a customer support agent. Increasingly, companies and organisations face challenges in addressing the various information needs of their customers, particularly given the growing diversity of user skills, knowledge levels or language preferences. There exists tremendous potential in this area for personalisation through the development of the proposed hybrid AH-PIR systems, which can potentially provide better assistance, guidance and navigation for individual users and their needs, preferences and context.

By making use of corporate knowledge encapsulated in professionally authored content, it is possible to mine models that include e.g. the various product features, their respective relationships, possible problems, etc. Moreover, by capturing a customer's interactions with such a portal, a hybrid system can estimate a user's expertise regarding a product or service and adapt its future result compositions accordingly. In addition, many other user characteristics can be taken into account, such as for example a user's input device, language preferences (possibly multilingual capabilities) or current query intent (e.g. learning about product features vs. reacting to a problem). These various user dimensions could all be used by a hybrid system in order to i) adapt a user's query (e.g. diversifying, narrowing), ii) adapt the retrieval (e.g. include more introductory explanations for a novice) and choose from various information sources (e.g. corporate website, forums), and iii) personalise the result composition (e.g. grouping, inter-linking, annotation, navigational cues) and presentation (e.g. according to the user's input device or modality preferences).

As mentioned in the previous section, it is crucial to consider the application and user task context when combining the various adaptation capabilities of a system. In the customer support scenario, an adaptive system could similarly be configured in order to suit particular situations.

First of all, if a user's current query intent is to find out about product functionalities, a broad-type navigation might be employed in order to provide an overview of the various features (as well as how these features relate to each other). In this case, broadening the initial query might be more advisable than narrowing the query. This would also allow the final composition and presentation step to provide a wider range of search results.

A second situation could be that a user is querying about a particular problem regarding the product (e.g. an error message that appeared while using the product). In this scenario, the user query could be enhanced by making use of the current work context (e.g. which feature of the product is currently being used) in order to narrow the query towards the specific problem. Also, retrieval adaptation might be employed by searching across specific information sources (e.g. community forums).

Thirdly, depending on the amount of prior experience of the user with the product, the query adaptation, retrieval adaptation and result composition might present this information to better support different information requirements. A novice user might benefit from the retrieval also being broadened towards more introductory information in order to fill a prerequisite knowledge gap. Similarly, this would enable the result composition to better provide an adaptive navigation across the various degrees of information difficulty (i.e. level of detail). On the other hand, this broadening of the original query might be regarded as unnecessary by an advanced user, who may prefer the query to be as precise as possible, regardless of the level of difficulty of the retrieved information.

As can be seen from these adaptation possibilities in a customer care information portal, a hybrid system is able to bridge the gap between different information sources, as well as varying user information needs and contexts. This is achieved by adaptively selecting and combining appropriate personalisation techniques depending on the user and application contexts. In the given scenario, this could prove beneficial not only for product and service consumers in terms of support satisfaction, but also for information providers in

terms of content reusability, customer support efficiency as well as customer loyalty. Moreover, since user queries are handled on an individual basis, customers can be motivated to search and consequently browse corporate websites for longer. This positive effect also provides possibilities in many other application areas, such as enabling e-learning systems to incorporate and personalise open-web information and consequently motivate users to read and learn more.

As mentioned in section 2, such a hybrid paradigm could also be considered a potential direction to take in *Exploratory Search*, as it can provide additional assistance for *open-ended information needs, in the context of specific user tasks and through a combination of browsing and focused searching*. Moreover, due to the historical educational background of AH systems, many related evaluation techniques can be used for *advanced system evaluation with respect to learning, insight, task outcomes and system utility*.

8. Future Directions

In order to achieve the desired goal of multidimensional open-web personalisation, there remain many challenges as well as opportunities for future research directions.

As identified in previous sections, the first challenge lies in deciding which part of the adaptation process to tailor for a particular situation. The trade-offs and interdependencies of adaptation techniques have to be understood and managed, especially considering that different application and user task contexts benefit variably from different adaptations. Also, depending on the available user information, it might be easier to adapt different parts of the process. For example, if a system only has access to very little information about the user (e.g. a few previous page visits), or if the system is not confident about the information regarding that user, it might be advisable to just adapt the user query. Similarly, if only prior knowledge information is available to the application, it might be more advisable to not adapt the user query, but rather to provide an adaptive navigation to guide the user through the retrieved information. Another key challenge therefore lies in building flexible system frameworks that can be easily configured for different application scenarios and varying user information availability. Such architectures would also allow the creation of overarching strategies across the search process in order to best link techniques and to avoid possible detrimental effects of adaptation interdependencies.

Since many of the metadata-based systems rely on sophisticated conceptual models, another challenge lies in the identification and generation of such models for a given application scenario. In particular, in order to apply adaptation and personalisation, the capturing and modelling of user preferences, interests and context is of paramount importance. Moreover, open-corpus user modelling can be characterised as particularly challenging, given the scattered availability of rich metadata on the web. Recent advances in named entity recognition as well as lightweight semantic model extraction show promising directions towards the generation of richer, open-corpus user and domain models (e.g. Sah and Wade (2011) propose metadata generation using fuzzy inferencing). However, it is also important to note again that *a little bit of personalisation goes a long way* and that adaptation and personalisation designers therefore only need to capture the *relevant* set of user attributes for a given application scenario. There is often no need for the generation of exhaustive user model attributes, and only a selection of the presented AH/PIR techniques needs to be applied at a given time.

As mentioned previously, privacy is another pressing issue and has been identified as one of the major challenges of personal services (Ashman et al., 2009). Data confidentiality and ownership are of major concern to many users, who might feel that systems gain too much insight into private aspects of their daily lives. Moreover, there is also the *accuracy* of any data that is inferred - personalisation needs to be sufficiently sophisticated to avoid "jumping to conclusions" based on inadequate information about a person, and to deal with exceptions when they occur. However, it is important to reiterate that most systems only need to have limited knowledge about their users to be able to personalise their web experiences. Moreover, it is recommended that systems incorporate the notions of transparency and scrutability in order to accurately show the knowledge that is currently held about them and/or any assumptions the system might have inferred as a consequence. Similarly, it is important to note that users should be in control of their user profiles and that they can see the clear benefits of having the personalisation enabled.

Finally, there are many social aspects that can potentially contribute to the personalisation experience of open web systems. In addition to traditional collaborative recommender systems, social applications such as Facebook¹⁰ and Twitter¹¹ have opened new possibilities for the development of much richer user/group modelling techniques. In particular, Facebook's Graph API¹², as well as Google's Social Graph API¹³ now allow the leveraging of not only a user's personal preferences, but also his/her social environment such as friends, family or work colleagues. Again, not all of this information may be needed for every personalised application scenario, but the careful selection and possible augmentation of such information can lead to many new application opportunities.

Acknowledgements

This research is supported by the Science Foundation Ireland (Grant 07/CE/I1142) as part of the Centre for Next Generation Localisation (<http://www.cngl.ie>) at Trinity College Dublin.

References

- Abel, F., Gao, Q., Houben, G.J., and Tao, K. (2011). Semantic Enrichment of Twitter Posts for User Profile Construction. *Proceedings of 8th Extended Semantic Web Conference (ESWC)*, pp. 375-389.
- Ahn, J.-W., Brusilovsky, B., Grady, J., He, D. and Florian, R. (2010). Semantic annotation based exploratory search for information analysts. *Information Processing and Management*, 46, pp. 383-402.
- Al-Maskaria, A., and Sanderson, M. (2011). The effect of user characteristics on search effectiveness in information retrieval. *Information Processing and Management*, 47, 5, pp. 719-729.
- Aroyo, L., De Bra, P., and Houben, G.J. (2003). Embedding Information Retrieval in Adaptive Hypermedia: IR meets AHA!. *Proceedings of the Workshop on Adaptive Hypermedia and Adaptive Web-Based Systems, at the 12th World Wide Web Conference, WWW2003*, pp. 63-76.
- Ashman, H., Brailsford, T. and Brusilovsky, P. (2009). Personal Services: Debating the Wisdom of Personalisation. In *Proceedings of the 8th International Conference on Advances in Web Based Learning (ICWL '09)*, pp. 1-11.
- Bakalov, F., König-Ries, B., Nauertz, A., and Welsch, M. (2010) Introspectiveviews: An interface for scrutinizing semantic user models. *Proceeding of the 18th International Conference on User Modeling, Adaptation, and Personalization, UMAP 2010*, pp. 219-230.
- BBC News. (2010). EU launches antitrust probe into alleged Google abuses. <http://www.bbc.co.uk/news/technology-11876443>, 30/11/2010.
- Berners-Lee, T., Hendler, J., and Lassila O. (2001). The Semantic Web. *Scientific American*, 284, 5, pp.35-43.
- Bhavnani, S., Bichakjian, C., Johnson, T., Little, R., Peck, F., Schwartz, J., and Strecher, V. (2003) Strategy hubs: next-generation domain portals with search procedures. *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 393-400.
- Bhogal, J., Macfarlane, A., and Smith, P. (2007). A review of ontology based query expansion. *Information Processing and Management*. 43, 4, pp. 866-886.
- Billsus, D., and Pazzani, M. J. (2007). Adaptive News Access. *The Adaptive Web*, LNCS, vol. 4321, pp. 550-570.
- Brin, S., and Page, L. (1998). The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30, 1-7, pp. 107-117.

¹⁰ <http://www.facebook.com>

¹¹ <http://www.twitter.com>

¹² <http://developers.facebook.com/docs/reference/api>

¹³ <http://code.google.com/apis/socialgraph>

- Brown, P. J., and Jones, G. J. F. (2001). Context-aware Retrieval: Exploring a New Environment for Information Retrieval and Information Filtering. *Personal Ubiquitous Computing*, 5, 4, pp. 253-263.
- Brusilovsky, P., and Henze, N. (2007). Open Corpus Adaptive Educational Hypermedia. *The Adaptive Web*, LNCS, vol. 4321, pp. 671-696.
- Brusilovsky, P., Chavan, G., and Farzan, R. (2004). Social adaptive navigation support for open corpus electronic textbooks. *Adaptive Hypermedia and Adaptive Web-Based Systems, AH 2004*, LNCS, vol. 3137, pp. 24–33.
- Burke, R. (2007). Hybrid web recommender systems. In *The Adaptive Web: Methods and Strategies of Web Personalization*, P. Brusilovsky, A. Kobsa, and W. Nejdl, Eds., LNCS, vol. 4321, pp. 377-408.
- Burton-Jones, A., Storey, V., Sugumaran, V., and Sandeep, P. (2003). A heuristic-based methodology for semantic augmentation of user queries on the Web. *Conceptual Modeling - ER 2003*, LNCS, vol. 2813, pp. 476-489.
- Calegari, S., and Pasi, G. Personalized Ontology-Based Query Expansion. (2008). *IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology*, vol. 3, pp. 256-259.
- Cantador, I., Fernández, M., Vallet, D., Castells, P., Picault, J., and Ribière, M. (2008). A Multi-Purpose Ontology-Based Approach for Personalised Content Filtering and Retrieval. *Advances in Semantic Media Adaptation and Personalization. Springer-Verlag, Studies in Computational Intelligence*, vol. 93, pp. 25-51.
- Chen, G., Bai, H., Shou, L., Chen, K., and Gao, Y. (2011). UPS: efficient privacy protection in personalized web search. *Proceedings of the 34th international ACM SIGIR Conference on Research and development in Information, SIGIR '11*, pp. 615-624.
- Chirita, P., Firan, C. S., and Nejdl, W. Personalized query expansion for the web. (2007). *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '07*, pp. 7-14.
- Conlan, O., Wade, V., Bruen, C. and Gargan, M. (2002). Multi-model, Metadata Driven Approach to Adaptive Hypermedia Services for Personalized eLearning. *Adaptive Hypermedia and Adaptive Web-Based Systems, AH2002*, LNCS, vol. 2347, pp. 100-111.
- Conlan, O., and Wade, V. (2004). Evaluation of APeLS – An Adaptive eLearning Service Based on the Multi-model, Metadata-Driven Approach. *Adaptive Hypermedia and Adaptive Web-Based Systems, AH2004*, LNCS, vol. 3137, pp. 291-295.
- Croft, B., Metzler, D. and Strohman, T. (2009). Search Engines: Information Retrieval in Practice (1st ed.). *Addison-Wesley Publishing Company*, USA.
- Daoud, M., Tamine, L., and Boughanem, M. (2010). A personalized graph-based document ranking model using a semantic user profile. *Proceedings of the 18th international conference on User modeling, adaption, and personalization, UMAP 2010*, pp. 171–182.
- De Bra, P., Aerts, A., Berden, B., de Lange, B., Rousseau, B., Santic, T., Smits, D., and Stash, N. (2003). AHA! The adaptive hypermedia architecture. *Proceedings of the Fourteenth ACM Conference on Hypertext and Hypermedia, HYPERTEXT '03*, pp. 81-84.
- Dey, A. (2001). Understanding and Using Context. *Personal and Ubiquitous Computing*, 5, 1, pp. 4–7.
- Dolog, P., Henze, N., Nejdl, W. and Sintek, M. (2003). Towards the adaptive Semantic Web. *Principles and Practice of Semantic Web Reasoning*, LNCS, vol. 2901, pp. 51-68.
- Dumais, S. Thinking Outside the (Search) Box. (2009). *Keynote at First International Conference on User Modeling, Adaptation, and Personalization, UMAP2009*.
- Farrell, R. G., Liburd, S. D., and Thomas, J. C. (2004). Dynamic assembly of learning objects. *Proceedings of the 13th international World Wide Web Conference on Alternate Track Papers & Posters, WWW Alt. '04*, pp. 162-169.

- Fernandez, M., Lopez, V., Sabou, M., Uren, V., Vallet, D., Motta, E., and Castells, P. (2008). Semantic Search Meets the Web. *Proceedings of the 2008 IEEE international Conference on Semantic Computing*, pp. 253-260.
- Geurts, J., Bocconi, S., and Van Ossenbruggen, J. (2003). Towards ontology-driven discourse: From semantic graphs to multimedia presentations. *The Semantic Web - ISWC 2003*, LNCS, vol. 2870, pp. 597-612.
- Glover, E. J., Lawrence, S., Gordon, M. D., Birmingham, W. P., and Giles, C. L. (2001). Web Search-Your Way. *Communications of the ACM*, 44, 12, pp. 97-102.
- Guha, R., McCool, R., and Miller, E. Semantic search. (2003). *Proceedings of the 12th International Conference on World Wide Web, WWW '03*, pp. 700-709.
- Haveliwala, T. H. (2003). Topic-Sensitive PageRank: A Context-Sensitive Ranking Algorithm for Web Search. *IEEE Transactions on Knowledge and Data Engineering*, 15, 4, pp. 784-796.
- Henze, N. (2005). Personal Readers: Personalized Learning Object Readers for the Semantic Web. *Proceedings of 12th International Conference on Artificial Intelligence in Education, AIED '2005*, pp. 274-281.
- Henze, N., and Nejd, W. (2001). Adaptation in open corpus hypermedia. *International Journal of Artificial Intelligence in Education*, 12, 4, pp. 325-350.
- Hirschman, L., and Gaizauskas, R. (2001). Natural language question answering: the view from here. *Natural Language Engineering*, 7, 4, pp. 275-300.
- Hsiao, I.-H., Brusilovsky, P., Yudelson, M., and Ortigosa, A. (2010). The value of adaptive link annotation in e-learning: a study of a portal-based approach. *Proceedings of the 21st ACM conference on Hypertext and hypermedia, HT '10*, pp. 223-228.
- Hsiao, I.-H., Sosnovsky, S., and Brusilovsky, P. (2009). Adaptive Navigation Support for Parameterized Questions in Object-Oriented Programming. *4th European Conference on Technology Enhanced Learning, EC-TEL 2009*, pp. 88-98.
- Jansen, B. J., Spink, A., and Saracevic, T. (2000). Real life, real users, and real needs: a study and analysis of user queries on the web. *Information Processing and Management*. 36, 2, pp. 207-227.
- Jansen, B. J., Booth, D. L., and Spink, A. (2008). Determining the informational, navigational, and transactional intent of Web queries. *Information Processing and Management*. 44, 3, pp. 1251-1266.
- Joachims, T., Granka, L., Pan, B., Hembrooke, H., and Gay, G. (2005). Accurately interpreting clickthrough data as implicit feedback. *Proceedings of the 28th international ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '05*, pp. 154-161.
- Joho, H., and Jose, J. M. (2008). Effectiveness of additional representations for the search result presentation on the web. *Information Processing and Management*, 44, 1, pp. 226-241.
- Jovanović, J., Gašević, D., and Devedžić, V. (2006). Dynamic Assembly of Personalized Learning Content on the Semantic Web. *The Semantic Web: Research and Applications*, LNCS, vol. 4011, pp. 545-559.
- Karam, N., Linckels, S. and Meinel, C. (2007). Semantic composition of lecture subparts for a personalized e-learning. *The Semantic Web: Research and Applications*, LNCS, vol. 4519, pp. 716-728.
- Kim, M., and Choi, K. (1999). A comparison of collocation-based similarity measures in query expansion. *Information Processing and Management*, 35, 1, pp. 19-30.
- Kleinberg, J. M. (1999). Authoritative sources in a hyperlinked environment. *Journal of the ACM*, 46, 5, pp. 604-632.
- Koutrika, G., Ioannidis, Y. (2005). A Unified User Profile Framework for Query Disambiguation and Personalization. *Proceedings of the Workshop on New Technologies for Personalized Information Access, PIA2005*, pp. 44-53.

- Krovetz, R., (1997). Homonymy and Polysemy in Information Retrieval. *Proceedings of the Eighth Conference on European Chapter of the Association for Computational Linguistics*, pp. 72–79.
- Kumaran, G. and Allan, J. (2008). Effective and efficient user interaction for long queries. *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '08*, pp. 11-18.
- Lin, Y.-l., and Brusilovsky, P. (2011). Towards open corpus adaptive hypermedia: a study of novelty detection approaches. *Proceedings of the 19th international conference on User modeling, adaption, and personalization, UMAP'11*, pp. 353-358.
- Linckels, S., Sack, H., and Meinel, C. (2007). Optimizing the retrieval of pertinent answers for nl questions with the e-librarian service. *Service Matchmaking and Resource Retrieval in the Semantic Web, SMR2*, pp. 139-152.
- Liu, F., Yu, C., and Meng, W. (2004). Personalized Web Search For Improving Retrieval Effectiveness. *IEEE Transactions on Knowledge and Data Engineering*, 16, 1, pp. 28–40.
- Mangold, C. (2007). A survey and classification of semantic search approaches. *International Journal of Metadata, Semantics and Ontologies*, 2, 1, pp. 23-34.
- Marchionini, G. (2006). Exploratory search: from finding to understanding. *Communications of the ACM*, 49, 4, pp. 41-46.
- Maekawa, T., Yanagisawa, Y., Sakurai, Y., Kishino, Y., Kamei, K. and Okadome, T. (2009) Web searching for daily living. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval (SIGIR '09)*, pp. 27-34.
- Micarelli, A., Gasparetti, F., and Biancalana, C. (2006). Intelligent search on the Internet. *Reasoning, Action and Interaction in AI Theories and Systems*, pp. 247–264.
- Micarelli, A., and Sciarrone F. (2004). Anatomy and empirical evaluation of an adaptive web-based information filtering system. *User Modeling and User-Adapted Interaction*, 14, 2-3, pp. 159-200.
- Micarelli, A., Gasparetti, F., Sciarrone, F., and Gauch, S. (2007). Personalized Search on the World Wide Web. *The Adaptive Web*, LNCS, vol. 4321, pp. 195-230.
- Miller, G. (1995). WordNet: a lexical database for English. *Communications of the ACM*, 38, 11, pp. 39-41.
- Navigli, R., and Velardi, P. (2003). An analysis of ontology-based query expansion strategies. *Workshop on adaptive text extraction and mining, ATEM 2003, at the 14th European conference on machine learning, ECML 2003*, pp. 42-49.
- Noh, H.-Y., Lee, J.-H., Oh, S.-W., Hwang, K.-S., and Cho, S.-B. (2011). Exploiting indoor location and mobile information for context-awareness service. *Information Processing and Management*, 48, 1, pp. 1-12.
- Pariser, E. (2011). *The filter bubble: What the Internet is hiding from you*. New York, NY: Penguin
- Pazzani, M. J. and Billsus, D. (2007). Content-based recommendation systems. In *The Adaptive Web: Methods and Strategies of Web Personalization*, P. Brusilovsky, A. Kobsa, and W. Nejdl, Eds., LNCS, vol. 4321, pp. 325-341.
- Pitkow, J., Schütze, H., Cass, T., Cooley, R., Turnbull, D., Edmonds, A., Adar, E., and Breuel, T. (2002). Personalized search. *Communications of the ACM*, 45, 9, pp. 50-55.
- Psarras, I., and Jose, J. (2006). A System for Adaptive Information Retrieval. *Adaptive Hypermedia and Adaptive Web-Based Systems, AH2006*, LNCS, vol. 4018, pp. 313-317.
- Radlinski, F. and Dumais, S. (2006). Improving personalized web search using result diversification. *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '06*, pp. 691-692.
- Rocchio, J. J. (1971). Relevance feedback in information retrieval. In G. Salton (Ed.), *The SMART retrieval system. Experiments in automatic document processing*, pp. 313-323. Englewood Cliffs, NJ: Prentice-Hall.

- Rocha, C., Schwabe, D., and Aragao, M. P. (2004). A hybrid approach for searching in the Semantic Web. *Proceedings of the 13th international Conference on World Wide Web, WWW '04*, pp. 374-383.
- Sah, M., and Wade, V. (2011). Automatic mining of cognitive metadata using fuzzy inference. *Proceedings of the 22nd ACM conference on Hypertext and hypermedia, HT '11*, pp. 37-46.
- Schafer, J. B., Frankowski, D., Herlocker, J., and Sen, S. (2007). Collaborative filtering recommender systems. In *The Adaptive Web: Methods and Strategies of Web Personalization*, P. Brusilovsky, A. Kobsa, and W. Nejdl, Eds., LNCS, vol. 4321, pp. 291-324.
- Sherman, C. (2011). Study Asks, Can You Trust Google's Personalized Search Results? <http://searchengineland.com/study-asks-can-you-trust-googles-personalized-search-results-64709>, 14/02/2011.
- Simonin, J., and Carbonell, N. (2008). Proactive Versus Multimodal Online Help: An Empirical Study. *Adaptive Hypermedia and Adaptive Web-Based Systems, AH2008*, LNCS, vol. 5149, pp. 183-192.
- Sintek, M, and S Decker. (2002). TRIPLE-A Query, Inference, and Transformation Language for the Semantic Web. *The Semantic Web - ISWC 2002*, LNCS, vol. 2342, pp. 364-378.
- Smith, A.S.G., and Blandford, A. (2003). MLTutor: An application of machine learning algorithms for an adaptive Web-based information system. *International Journal of Artificial Intelligence in Education*, 13, 2-4, pp. 235-261.
- Smits, D., and De Bra, P. (2011). GALE: a highly extensible adaptive hypermedia engine. *Proceedings of the 22nd ACM conference on Hypertext and hypermedia, HT '11*, pp. 63-72.
- Speretta, M., and Gauch, S. (2005). Personalized Search Based on User Search Histories. *Web Intelligence, WI2005*, pp. 622-628.
- Spink, A., and Jansen, B.J. (2004). A study of web search trends. *Webology* 1, 2. <http://www.webology.ir/2004/v1n2/a4.html>.
- Stamou, S., and Ntoulas, A. (2009). Search personalization through query and page topical analysis. *User Modeling and User-Adapted Interaction*, 19, 1-2, pp. 5-33.
- Sushmita, S., Joho, H., Lalmas, M. and Villa, R. (2010). Factors affecting click-through behavior in aggregated search interfaces. *Proceedings of the 19th ACM international conference on Information and knowledge management, CIKM '10*, pp. 519-528.
- Tanudjaja, F., and Mui, L. (2002). Persona: A Contextualized and Personalized Web Search. *Proceedings of the 35th Annual Hawaii international Conference on System Sciences, HICSS '02*, 3, 3, pp. 67.
- Teevan, J., Dumais, S. T., and Horvitz, E. (2005). Personalizing search via automated analysis of interests and activities. *Proceedings of the 28th international ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '05*, pp. 449-456.
- Teevan, J., Dumais, S., and Gutt, Z. (2008). Challenges for supporting faceted search in large, heterogeneous corpora like the Web. *Proceedings of HCIR 2008*, pp. 6-8.
- Thomas, P., Noack, K., and Paris, C. (2010). Evaluating interfaces for government metasearch. In *Proceeding of the third symposium on Information interaction in context (IIiX '10)*, pp. 65-74.
- Tran, T., Wang, H., Lamparter, S., and Cimiano, P. (2008). Personalization Using Ontologies and Rules. *Adaptive Hypermedia and Adaptive Web-Based Systems, AH2008*, LNCS, vol. 5149, pp. 349-352.
- Truran, M., Goulding, J., and Ashman, H. (2005). Co-active intelligence for image retrieval. *Proceedings of the 13th Annual ACM international Conference on Multimedia*, pp. 547-550.
- Tsandilas, T. and schraefel, m.c. (2004). Usable adaptive hypermedia systems. *New Review of Hypermedia and Multimedia*, 10, 1, pp. 5-29.
- Tvarozek, M. and Bielikova, M. (2007). Adaptive faceted browser for navigation in open information spaces. *Proceedings of the 16th international Conference on World Wide Web, WWW '07*, pp. 1311-1312.

- Villa, R., Cantador, I., Joho, H., and Jose, J. M. (2009). An aspectual interface for supporting complex search tasks. *In Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval (SIGIR '09)*, pp. 379-386.
- Wade, V. Challenges for the Multi-Dimensional Personalised Web. (2009). *Keynote at First International Conference on User Modeling, Adaptation, and Personalization, UMAP2009*.
- White, R. W., Jose, J. M., and Ruthven, I. (2003). A task-oriented study on the influencing effects of query-biased summarisation in web searching. *Information Processing and Management*, 39, 5, pp. 707-733.
- White, S. and Smyth, P. (2003). Algorithms for estimating relative importance in networks. *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '03*, pp. 266-275.
- White, R. W., and Roth, R. A. (2009). Exploratory Search: Beyond the Query-Response Paradigm, *Morgan and Claypool*.
- Xiang, B., Jiang, D., Pei, J., Sun, X., Chen, E., and Li, H. (2010). Context-aware ranking in web search. *Proceedings of the 33rd international ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '10*, pp. 451-458.
- Xu, J., and Croft, W. B. (1996). Query expansion using local and global document analysis. *Proceedings of the 19th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '96*, pp. 4-11.
- Xu, S., Jin, T., and Lau, F. C. (2009). A new visual search interface for web browsing. *Proceedings of the Second ACM international Conference on Web Search and Data Mining*, pp. 152-161.
- Yee, K.P., Swearingen, K., Li, K. and Hearst, M. (2003). Faceted metadata for image search and browsing. *In Proceedings of the SIGCHI conference on Human factors in computing systems (CHI '03)*, pp. 401-408.
- Zhang, L., and Zhang, Y. (2010). Interactive retrieval based on faceted feedback. *Proceedings of the 33rd international ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '10*, pp. 363-370.

Vitae

Ben Steichen holds a B.Sc. and an M.Sc. in Computer Science and is currently completing his Ph.D. in the School of Computer Science and Statistics at Trinity College Dublin. His research interests are in Adaptive Web and Personalisation, focusing on applying the adaptivity of Adaptive Hypermedia systems across the content scale of traditional Information Retrieval systems in order to achieve Open Adaptive Hypermedia.

Helen Ashman has worked in hypertext and Web research for many years, with interests in hypertext modelling, using click data, adaptive Web and user modelling, behavioural intrusion detection, among others. She is Editor-in-Chief of ACM Transactions on the Web and is on the Editorial Boards of the World Wide Web Journal and the New Review of Multimedia and Hypermedia.

Prof. **Vincent Wade** is Head of the Intelligent Systems Discipline in the School of Computer Science and Statistics in TCD. He is also Deputy Director of Centre for Next Generation Localisation (CNGL). Vincent holds BSc, MSc and PhD degrees in computer science. His research focuses on Adaptive Web systems and Personalisation and he has published in excess of two hundred technical papers in journal and international conferences of repute. He is a guest editor and reviewer for many IEEE, ACM journals and conferences including IEEE Communications, ACMs Transaction on the Web, Adaptive Hypermedia (now called User Modelling, Adaptation & Personalisation).